

PCT

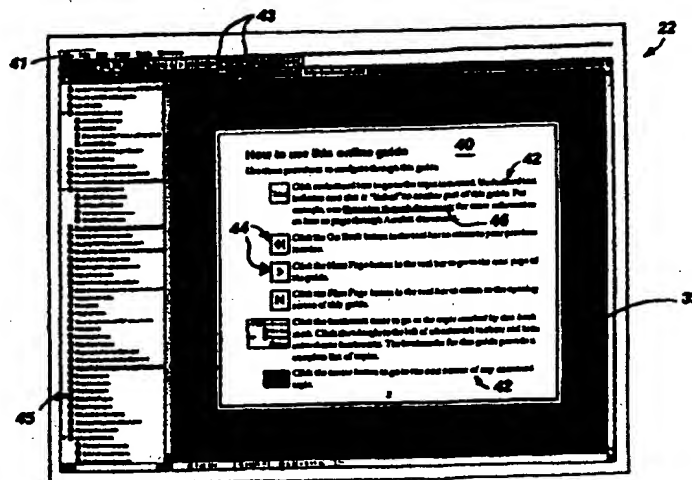
WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 15/16, 17/21, 13/14		A1	(11) International Publication Number: WO 97/12328
			(43) International Publication Date: 3 April 1997 (03.04.97)
(21) International Application Number: PCT/US96/15725		(81) Designated States: CA, JP, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).	
(22) International Filing Date: 25 September 1996 (25.09.96)		Published With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.	
(30) Priority Data: 08/533,177 25 September 1995 (25.09.95) US 08/533,875 26 September 1995 (26.09.95) US 08/569,000 7 December 1995 (07.12.95) US			
(71) Applicant: ADOBE SYSTEMS INCORPORATED [US/US]; 345 Park Avenue, San Jose, CA 95110 (US).			
(72) Inventors: ROWE, Edward, R.; 701 West 32nd Street #14, Los Angeles, CA 90007 (US). PRIYADARSHAN, Eswar; 1054 Heatherston Avenue, Sunnyvale, CA 94087 (US). ANDERSON, Kenneth, S.; 4133 Amaranta, Palo Alto, CA 94306 (US). AL-SHAMMA, Nabelle, A.; 1525 Gretel Lane, Mountain View, CA 94040 (US). TAFT, Edward, A.; 1147 Sladky Avenue, Mountain View, CA 94040 (US). MCQUARRIE, Elizabeth, M.; 15882 Ravine Road, Los Gatos, CA 95030 (US). COHN, Richard, J.; 575 North California Avenue, Palo Alto, CA 94301 (US).			
(74) Agent: TROESCH, Hans, R.; Fish & Richardson P.C., Suite 100, 2200 Sand Hill Road, Menlo Park, CA 94025 (US).			

(54) Title: OPTIMUM ACCESS TO ELECTRONIC DOCUMENTS



(57) Abstract

A method and apparatus for providing an optimized page-based electronic document file (figure 2a) and downloading the optimized file. An optimized document file is created from a non-optimized electronic document (74). Page contents are contiguously written in the optimized file and page offset information (68, fig. 8a) used to locate individual pages and objects of the document. When downloading the optimized file from a host, the page offset information is read early and is used to download a specific page requested by the user without downloading other pages in the document. A viewer may download (226) a first portion of the requested page, while all remaining portions of the requested page are located by a finder process using the page offset table (68, fig. 8a). Alternatively, all objects for a full page may be requested at once. The requested page can thus be downloaded with only one connection to the host. Optionally, shared objects can be downloaded and interleaved between portions of the page contents (284) that reference the shared objects. They can also be read in one transaction identifying byte ranges in the document. Subsequently, the requested page is displayed to the user on an output display device.

BEST AVAILABLE COPY

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AM	Armenia	GB	United Kingdom	MW	Malawi
AT	Austria	GE	Georgia	MX	Mexico
AU	Australia	GN	Guinea	NE	Niger
BB	Barbados	GR	Greece	NL	Netherlands
BE	Belgium	HU	Hungary	NO	Norway
BF	Burkina Faso	IE	Ireland	NZ	New Zealand
BG	Bulgaria	IT	Italy	PL	Poland
BJ	Benin	JP	Japan	PT	Portugal
BR	Brazil	KE	Kenya	RO	Romania
BY	Belarus	KG	Kyrgyzstan	RU	Russian Federation
CA	Canada	KP	Democratic People's Republic of Korea	SD	Sudan
CF	Central African Republic	KR	Republic of Korea	SE	Sweden
CG	Congo	KZ	Kazakhstan	SG	Singapore
CH	Switzerland	LI	Liechtenstein	SI	Slovenia
CI	Côte d'Ivoire	LK	Sri Lanka	SK	Slovakia
CM	Cameroon	LR	Liberia	SN	Senegal
CN	China	LT	Lithuania	SZ	Swaziland
CS	Czechoslovakia	LU	Luxembourg	TD	Chad
CZ	Czech Republic	LV	Latvia	TG	Togo
DE	Germany	MC	Monaco	TJ	Tajikistan
DK	Denmark	MD	Republic of Moldova	TT	Trinidad and Tobago
EE	Estonia	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	UG	Uganda
FI	Finland	MN	Mongolia	US	United States of America
FR	France	MR	Mauritania	UZ	Uzbekistan
GA	Gabon			VN	Viet Nam

OPTIMUM ACCESS TO ELECTRONIC DOCUMENTS**BACKGROUND OF THE INVENTION**

The present invention relates generally to the storage and retrieval of data for a computer system, and
5 more particularly to a method and apparatus for optimizing page-based data documents for fast retrieval over networks, and to a method and apparatus for accessing such optimized documents. The present invention also relates to methods and apparatus for the
10 processing and display of electronic documents, and more particularly to the processing and display of such documents when retrieved over networks.

It has become increasingly common to create, transmit, and display documents in electronic form.
15 Electronic documents have a number of advantages over paper documents including their ease of transmission, their compact storage, and their ability to be edited and/or electronically manipulated. An electronic document typically has information content (such as text,
20 graphics, and pictures) and formatting information that directs how the content is to be displayed. With recent advances in multimedia technology, documents can now also include sound, full motion video, and other multimedia content.

25 An electronic document is provided by an author, distributor or publisher (referred to as "publisher" herein) who often desires that the document be viewed with the appearance with which it was created. This, however, creates a problem in that electronic documents
30 are typically widely distributed and, therefore, can be viewed on a great variety of hardware and software platforms. For example, the video monitors being used to view the document can vary in size, resolution, etc. Furthermore, the various software platforms such as DOS,
35 Microsoft WindowsTM, and MacintoshTM all have their own

- 2 -

display idiosyncrasies. Also, each user or "reader" of the electronic document will have his or her own personal viewing preferences, which should be accommodated, if possible.

5 A solution to this problem is to provide a "portable electronic document" that can be viewed and manipulated on a variety of different platforms and can be presented in a predetermined format where the appearance of the document as viewed by a reader is as it
10 was intended by the publisher. One such predetermined format is the Portable Document Format™ (PDF™) developed by Adobe Systems, Inc. of Mountain View, California. An example of page-based software for creating, reading, and displaying PDF documents is the Acrobat™ software, also
15 of Adobe Systems, Inc. The Adobe Acrobat software is based on Adobe's PostScript® technology, which describes formatted pages of a document in a device-independent fashion. An Acrobat program on one platform can create, display, edit, print, annotate, etc. a PDF document
20 produced by another Acrobat program running on a different platform, regardless of the type of computer platform used. A document in a certain format or language can be translated into a PDF document using Acrobat. A PDF document can be quickly displayed on any
25 computer platform having the appearance intended by the publisher, allowing the publisher to control the final appearance of the document.

One relatively new application for portable electronic documents is the retrieval of such documents
30 from the "Internet", the globally-accessible network of computers that collectively provides a large amount and variety of information for users. From services of the Internet such as the World Wide Web, users may retrieve or "download" data from Internet network sites and
35 display the data that includes information presented as

- 3 -

text in various fonts, graphics, images, and the like having an appearance intended by the publisher. A file format such as PDF that allows any platform to view a document having an appearance as intended by a publisher
5 is thus of great value when downloading files from such widely-accessible and platform-independent network sources such as the Internet.

One problem with previous page-based data downloading processes is that all of the data of a
10 document is typically downloaded before any portion of the document is displayed to the user. Thus, the user must wait for an entire document to download before seeing a page or other portion of the document on the display screen. This can be inconvenient when the user
15 wishes to use only a portion of the document, i.e., view only specific pages or a specific number of contiguous pages of a document. Some searching processes allow a word to be searched in a document and will download only the portion of the document that includes the searched
20 word. However, this portion of the document is an isolated, separate portion that has no connection with the rest of the document. If the user wishes to view the next page after the downloaded portion, he or she must inconveniently either download the entire document or
25 specify a search term on the next page of the document.

Acrobat and similar programs for displaying portable electronic documents such as PDF documents are often page-based, which means that the program typically organizes and displays a desired page of the document at
30 a time. Typically, the entire document was downloaded at once, then desired pages displayed. However, Acrobat is conducive to downloading a page of a document at a time from a document file, while still allowing a user to select other pages of the document conveniently.
35 However, for such page-based formats, the document data

- 4 -

usually is not stored contiguously in a page order within a file, data structure, or other collection of document data ("document file" as referred to herein). For example, a document file in the PDF format may store a page having objects such as a page contents object (including text, graphics shapes, display instructions, etc.) and image objects. However, the objects may be stored in the document in a scattered or disjointed manner. For example, portions of the page contents object can be scattered in different places in a document file, and shared objects such as fonts can be stored anywhere in the file. Shared objects such as fonts can also be stored in files distinct from the document file, and even on a separate computer, or be made available through a resource service such as a font server. Since the output display device displays the page contents and shared objects based upon pointers to related objects, objects do not have to be stored sequentially or contiguously in the document file, and are typically stored in a disjointed manner.

This disjointed data storage for pages can lead to problems when attempting to download a specific page of a document desired by the user. One major problem is time delays caused by making multiple connections (or multiple request-response transactions) when downloading data. For example, a viewing program for displaying page-based data at a client computer begins downloading a PDF (or similar format) file from a remote host computer. The viewing program makes one connection to (or initiates one transaction with) the host and downloads data from the first portion of the page, then must make another connection to (or transactions with) the host to retrieve the next, disjointed portion of the page. This has the effect of slowing down the downloading of the page, since each connection (and each transaction) has a time delay

- 5 -

and overhead associated with it. The user requesting the page thus may have to wait several seconds before the viewer receives all of the data for the page and displays the page. This problem is compounded when fonts or other
5 such referenced objects are included on the page, since yet another connection must be made to (or transaction made with) the host to retrieve these objects before the page can be displayed.

The time delays for downloading a page can become
10 even lengthier when a randomly-accessed page is desired to be viewed by the user. In PDF files, objects are provided in a "page tree" which the viewer consults to determine where in the document file the root of a randomly-accessed page is positioned. The page tree is a
15 data structure in which every node must be visited in order to determine all the children objects in the tree. Thus, many page nodes may need to be visited to determine where a page root object is located in the document file. The page tree can thus be quite large, and downloading it
20 from the document slows the downloading process. In addition, the page tree is often so large or disjointed that multiple connections to (or transactions with) the host are required to download it.

Therefore, there is a need for a method and
25 apparatus for providing optimized page-based documents and downloading desired pages from such documents without causing an excessive delay before displaying a page, or portions of a page, to the user.

SUMMARY OF THE INVENTION

30 The present invention provides a method and apparatus for optimizing a page-based electronic document and downloading and displaying desired pages, or portions of a page, from the optimized document without excessive time delays.

- 6 -

A method of the present invention provides an optimized document file from a non-optimized electronic document having one or more pages. Page content information that describes individual pages of the document is written in the optimized document file. The page content information may be written contiguously. Page offset information used to locate individual pages of the document may also be provided in the optimized document file. Objects shared by multiple pages are also provided for in the optimized document file, contiguously located after all of the page content information, and the page offset information includes offsets (locations) to these shared objects. The page content information includes text and graphics, and the shared objects can include font objects and image objects. To provide the page contents and shared objects contiguously in the file, an internal list of non-shared objects and shared objects in the document file is created. A list of pages that share objects is also created that includes the shared objects used by each sharing page and, for each such shared object, a portion of the page contents in which the shared object is referenced. In addition, in one aspect, first page offset information may be provided in a range table for a first page of the optimized document file. Such first page offset information describes the locations of all portions of the first page in the document file. The offsets to page content for this page may be interleaved in the range table with offsets to shared objects referenced by the page content for the first page.

Another method of the present invention efficiently downloads a page-based optimized document created as described above. The page offset information is read early during the downloading process. Beginning and ending offsets of each page of the document can be

- 7 -

derived from the page offset information. Using the page offset information, a specific page requested by the user is downloaded, and any page desired by the user can readily be downloaded without the necessity of

5 downloading other pages in the document. In one aspect of the method, the page offset information may be read before the downloading of more than one page of the document has been completed. In one aspect, the document file has a pointer that points to the location of the

10 page offset information, which pointer is read ahead of, or during, the reading of the first page of the document.

In another aspect, when a user requests a specific page of an optimized document, the specific page is downloaded to a client computer system in only one

15 connection with a host that stores the optimized document file. In another aspect, the specific page is downloaded in only one transaction with the host. The requested page, while being downloaded, may be displayed to the user on an output display device, such as a display

20 screen, monitor, or printer. The downloading can be accomplished by a viewer program on the client computer system. When connecting and downloading, the viewer may download a first portion of the requested page, while all remaining portions of the requested page are located and

25 requested by a finder process on the client computer using the page offset table. These additional portions are downloaded during the client computer's one connection with the host, thus saving time and overhead by avoiding multiple transactions or connections. The

30 additional portions of the specific page may include shared objects referenced by page contents of the specific page. Shared objects are downloaded in an interleaved order between portions of the page contents that reference the shared objects. In another aspect,

35 the requested page is downloaded to a client computer

- 8 -

system in only one transaction with a host that stores the optimized document file, the transaction being constructed by a process using a page offset hint table and any other hint tables available in the document.

5 If shared objects are downloaded in an interleaved order, the interleaving process includes downloading a first portion of page content from the requested page, where the first portion of page content includes a reference to a shared object. The first portion may
10 include all contiguous page content of the document until the (approximate) point of reference to the shared object. Then, the shared object referenced by the first portion of the page is downloaded. The shared object is, for example, a font or similar referenced object that is
15 needed to display the first portion of page content. A second portion of page content from the requested page is then downloaded, where the second portion is contiguous with the first portion of page content. The locations of the first and second portions of page content and the
20 shared objects in the page-based document are derived using the page offset table. Alternatively, a surrogate, such as a substitute font, is used to display the first portion of page content, thereby allowing the process to defer the downloading of the referenced object and
25 thereby to download and to display more quickly the second portion of page content.

Another method of the present invention provides for the displaying on a display device of a computer an electronic document, such as a portable electronic
30 document, having text to be displayed on top of a large object, such as a bitmap image. In general, in one aspect, the method includes displaying the display of the large object in favor of displaying the overlying text, displaying the overlying text on the display device, and
35 at least as to that portion of the large object that

- 9 -

appears underneath the overlying text, drawing the underneath portion into an off-screen buffer, drawing the overlying text over the object in the off-screen buffer and copying the off-screen buffer to be displayed on the display device. In another aspect, the acts of displaying an object and of displaying text include rendering a bitmap of at least one bit per pixel into a display buffer of random access memory. In another aspect, the display buffer and the off-screen buffer have the same pixel depths and color definitions. In another aspect, the invention provides for creating a blocking mask corresponding to the displayed appearance of the text and then displaying the portion of the object that is specified to appear as if drawn underneath the text under control of the blocking mask so that displaying the object does not overwrite the displayed text.

Another method of the present invention is implemented in a viewing program to display to a user an electronic document, such as a portable electronic document, that contains an interactive element responsive to user input. In one aspect, the method includes changing the appearance of the cursor of the viewing program's graphical user interface to indicate when it is located in a position where the interactive element will be displayed, and making the interactive element responsive to input from the user without waiting for the interactive element to be displayed. In another aspect, the interactive element is a hypertext link. In another aspect, the interactive element is an annotation in a PDF format electronic document.

Another method of the present invention provides for displaying on a computer display device an electronic document, such as a portable electronic document, that has text in a desired font, without waiting for the desired font to be available. In one aspect, the method

- 10 -

includes initially drawing on the display device at least a part of the text in a substitute font different from the desired font, obtaining the desired font for use on the computer with the display device, and redrawing with
5 the desired font the area of display in which the substitute font had been used initially. In another aspect, the method includes reading font description metrics for the desired font and using them to create a substitute font. In another aspect, the method also
10 includes adopting a font from available font resources as the substitute font. In another aspect, the desired font is a font included as an embedded font in the document. In another aspect, the desired font is obtained from a font server.

15 An apparatus of the present invention provides for efficiently downloading a page-based document stored on a host, as described above. The apparatus includes a digital processor, a memory device, and a display screen. Furthermore, a mechanism for displaying the page-based
20 document on the display screen is included which connects with the host to download the page offset information and/or to download a specific page of the document requested by the user without downloading other pages in the document. A downloaded page can be displayed on the
25 display screen. A finder uses the page offset information to provide a location of the specific page in the document to the displaying mechanism so that the specific page can be downloaded. The finder can request additional portions of the specific page during one
30 connection, and can interleave the downloading of portions of page contents and shared objects. In another aspect, the finder can request all portions of the specific page during one transaction, the transaction being constructed by a process using a page offset hint
35 table and any other hint tables available in the

- 11 -

document.

In general, in one aspect, the invention provides a method for reading a user-requested page of an electronic document that contains page content elements and a set of hint tables, where one of the set of the hint tables (a page offset hint table) provides page offset information for pages of the document. The method includes reading the page offset hint table early during reading of the document, and using the page offset hint table to locate the contents of the user-requested page. In this way, the user-requested page can readily be read without the necessity of reading other pages in the document. In another aspect, the page offset information is read before the reading of more than one page of the document has been completed, and the beginning and ending offsets of each page of the document can be derived from the page offset information. In another aspect, the page offset information is read before the reading of a second page of the document is begun. In another aspect, the set of hint tables includes a shared object hint table that is read before the reading of a second page of the document is begun. In another aspect, the set of hint tables includes a bookmark hint table that is read before the reading of a second page of the document is begun. In another aspect, the set of hint tables includes an article thread hint table that is read before the reading of a second page of the document is begun. In another aspect, the set of hint tables includes a thumbnail hint table that is read before the reading of a second page of the document is begun. In another aspect, the document is read from a host computer by a viewer program running on a user computer and the invention includes displaying the specific page requested by the user on a display device coupled to the user computer. In another aspect, each category of objects associated with the document as

- 12 -

a whole has a corresponding hint table. In another aspect, the user-requested page is downloaded in one transaction with the host computer. In another aspect, the entire document is requested in an initial

5 transaction with the host computer, and the reading of the document from the host computer is not interrupted until the user requests a specific page of the document. In another aspect, shared objects are cached after they are read.

10 In general, in one aspect, the invention provides a method for providing an optimized electronic document having two or more pages. The method includes providing document information in the optimized electronic document including page content information specifying the
15 appearance of each page of the document, providing the page content information specifying any aspect of the appearance of a designated first page of the document at the beginning of the optimized electronic document without regard to the nature or amount of the page
20 content information for the designated first page, and providing a page offset hint table in the optimized electronic document having information sufficient to locate respective page content information for each of substantially all the pages of the document. In another
25 aspect, the method includes providing page content information including text, at least one graphic, at least one image, and at least one font. In another aspect, the method includes providing a set of hint tables at the beginning of the optimized electronic
30 document. In another aspect, the method includes providing a set of hint tables in the optimized electronic document before page content information other than the page content information specifying any aspect of the appearance of the designated first page of the
35 document.

- 13 -

In general, in one aspect, the invention provides a method for reading an electronic document containing a set of objects, each object belonging to one of a set of classes, the document also containing a hint table for each class in the set of classes, each hint table containing information identifying the locations of objects in the corresponding class. The method includes reading the hint tables early during the reading process, using a hint table to identify the locations in the document of objects of the corresponding class, and using the identified locations to create a transaction to read all or a specified subset of the objects of the corresponding class. In this way, any aspect of the document related to a class of information desired by the user can be read without the necessity of reading objects relating to other aspects of the document. In another aspect, the set of classes includes a class for page objects. In another aspect, the set of classes includes a class for thumbnail objects. In another aspect, the set of classes includes a class for article objects. In another aspect, the set of classes includes a class for bookmark objects. In another aspect, the set of classes includes a class defined by an application plug-in having a corresponding hint table.

An advantage of the present invention is that downloading can be limited to only specific, desired pages of a document located on a remote host, resulting in a faster display time for those pages since the entire document need not be downloaded. In addition, a page-based electronic document file is optimized to include contiguously-stored page contents and a page offset table. The speed of downloading individual pages from the optimized file is increased by determining page locations in the optimized file from the page offset table. All page contents and shared objects for a page

- 14 -

can be downloaded at one connection to the host, thus avoiding the time delays of multiple connections.

Another advantage of this invention is that portions of page contents can be downloaded in an interleaved order with shared objects such as fonts needed to display those portions of page contents. This allows a downloaded portion of the page to be displayed more quickly without having to wait for referenced shared objects to be downloaded.

Another advantage of this invention is that undesirable delays in the displaying of text are reduced and that useful aspects of a desired page are more quickly made available to the user.

These and other advantages of the present invention will become apparent to those skilled in the art upon a reading of the following specification of the invention and a study of the several figures of the drawing.

BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings, which are incorporated in, and constitute a part of, the specification, schematically illustrate specific embodiments of the invention and, together with the general description given above and the detailed description of the embodiments given below, serve to explain the principles of the invention.

Figure 1 is a block diagram of a computer system for providing an optimized document and/or downloading data from an optimized file in accordance with the present invention.

Figure 2a is a diagrammatic illustration of a display screen of a viewer displaying a page from a portable electronic document and a bookmark view.

Figure 2b is a diagrammatic illustration of the

- 15 -

display screen of the viewer displaying a page from a portable electronic document and a "thumbnails" window.

Figure 3a is a diagrammatic illustration of an example of a non-optimized document file.

5 Figure 3b is a diagrammatic illustration of an optimized document file of the present invention.

Figure 4 is a flow diagram illustrating the process of the present invention for creating an optimized document file from an electronic document.

10 Figure 5 is a flow diagram illustrating a step of Figure 4 in which an internal list of objects and lists of shared objects are created.

Figure 5a is a diagrammatic illustration of a shared object list created in the process of Figure 5.

15 Figure 6 is a flow diagram illustrating a step of Figure 5 in which shared objects are processed.

Figure 6a is a diagrammatic illustration of a sharing pages list created in the process of Figure 6.

20 Figure 7 is a flow diagram illustrating a step of Figure 5 in which the sharing pages list is completed and shared objects are incorporated into the internal list.

Figure 8 is a flow diagram illustrating a step of Figure 4 in which a page offset table of the present invention is placed in the optimized file.

25 Figure 8a is a diagrammatic illustration of a page offset table created in the process of Figure 8.

Figure 9 is a flow diagram illustrating a step of Figure 4 in which a range table is written into the optimized file.

30 Figure 9a is a diagrammatic illustration of the range table created in the process of Figure 9.

Figure 10 is a flow diagram illustrating a process of the present invention for downloading pages from an optimized document file.

- 16 -

Figure 11 is a flow diagram illustrating a step of Figure 10 in which a finder retrieves the page offset table from the document and determines page information.

Figure 11a is a diagrammatic illustration of page 5 information tables developed in the process of Figure 11.

Figure 12 is a flow diagram illustrating a step of Figure 10 in which the finder requests any additional ranges of page data for a requested page during a viewer's connection with a host.

10 Figure 13a is a flow diagram illustrating an optional process of the present invention in a step of Figure 10 in which text is progressively rendered using different fonts.

Figure 13b is a flow diagram illustrating an 15 optional process of the present invention in a step of Figure 10 in which an active element of a document is enabled early.

Figure 13c is a flow diagram illustrating an optional process of the present invention in a step of 20 Figure 10 in which text is displayed before an underlying object is displayed.

Figure 14 is a block diagram illustrating a linearized document file of the present invention.

Figure 15a is a diagrammatical illustration of a 25 page offset hint table.

Figure 15b is a diagrammatical illustration of a shared object hint table.

Figure 15c is a diagrammatical illustration of a thumbnail hint table.

30 Figure 15d is a diagrammatical illustration of a generic object hint table.

- 17 -

DETAILED DESCRIPTION

The present invention is well-suited for downloading pages of data of a portable electronic document from a host computer, and for their optimum, incremental display. More particularly, a portable electronic document in Portable Document Format (PDF) and like formats can be optimized by the present invention. The present invention is suitable for organizing and downloading page-based files such that the time for downloading and displaying pages of the file is minimized.

In Figure 1, a computer system 10 for downloading a portable electronic document can include a digital computer 11, a display screen 22, a printer 24, a floppy disk drive 26, a hard disk drive 28, a network interface 30, and a keyboard 34. Other types of peripherals can also be included, such as a CD-ROM drive, input tablet or other interface devices, etc. Digital computer 11 typically includes a microprocessor 12, a memory bus 14, random access memory (RAM) 16, read only memory (ROM) 18, a peripheral bus 20, and a keyboard controller 32. Digital computer 11 can be a personal computer (such as an IBM-PC AT-compatible or Macintosh personal computer), a workstation (such as a SUN or Hewlett-Packard workstation), etc.

Microprocessor 12 is a general purpose digital processor which controls the operation of computer system 10. Microprocessor 12 can be a single-chip processor or can be implemented with multiple components. Using instructions retrieved from memory, microprocessor 12 controls the reception and manipulation of input data and the output and display of data on output devices. In the described embodiment, a function of microprocessor 12 is to read and process data from pages of an electronic document.

- 18 -

Memory bus 14 is used by microprocessor 12 to access RAM 16 and ROM 18. RAM 16 is used by microprocessor 12 as a general storage area and as scratch-pad memory, and can also be used to store
5 downloaded data that is being displayed (or not being displayed). ROM 18 can be used to store instructions followed by microprocessor 12 and other permanent data.

Peripheral bus 20 is used to access the input, output, and storage devices used by digital computer 11.
10 In the described embodiment, these devices include display screen 22, printer device 24, floppy disk drive 26, hard disk drive 28, and network interface 30. Keyboard controller 32 is used to receive input from keyboard 34 and send decoded symbols for each pressed key
15 to microprocessor 12 over bus 33.

Display screen 22 is an output device that displays images of data provided by microprocessor 12 via peripheral bus 20 or provided by other components in the computer system. In the described embodiment, display
20 screen 22 is a raster device which displays images on a screen corresponding to bits of a bitmap in rows and columns of pixel, as is well known to those skilled in the art. That is, a bitmap can be input to the display screen 22 and the bits of the bitmap can be displayed as
25 pixels. An input bitmap can be directly displayed on the display screen, or components of computer system 10 can first render codes or other image descriptions from a page description file into bitmaps and send those bitmaps to be displayed on display screen 24, as is also well
30 known. Raster display screens such as CRT's, LCD displays, etc. are suitable for the present invention.

Printer device 24 provides an image of a bitmap on a sheet of paper or a similar surface. Printer 24 can be a laser printer, which, like display screen 22, is a
35 raster device that displays pixels derived from bitmaps.

- 19 -

Printer device 24 can print images derived from data such as found in a portable electronic document. Other output devices can be used as printer device 24, such as a plotter, typesetter, etc. Computer system 10 can display
5 images on a display output device, such as display screen 22 or printer 24, using data from memory, a storage device, or from another source or host over a network connected by network interface 30.

Floppy disk drive 26 and hard disk drive 28 can be
10 used to store data such as a document that has been downloaded or created in the optimized format of the present invention. Floppy disk drive 26 facilitates transporting such data to other computer systems, and hard disk drive 28 permits fast access to large amounts
15 of stored data. Other mass storage units such as nonvolatile memory (e.g., flash memory), PC-data cards, or the like, can also be used to store data used by computer system 10. Herein, a "computer (or machine) readable storage medium" can refer to both memory such as
20 RAM 16 and ROM 18 as well as disk drives 26 and 28 or any other type of device for storing data.

Network interface 30 is used to send and receive data over a network connected to one or more other computer systems, such as computer device 31. An
25 interface card, modem, or similar device and appropriate software implemented by microprocessor 12 can be used to connect computer system 10 to an existing network and transfer data according to standard protocols. In the present invention, network interface 30 can be used to
30 retrieve or "download" portable electronic documents from a host computer system over a network, or send ("upload") the documents to a host or client computer system. The network can be implemented using a variety of hardware and software, as is well known to those skilled in the
35 art.

- 20 -

Keyboard 34 is used by a user to input commands and other instructions to computer system 10. Images displayed on display screen 22 or accessible to computer system 10 can be edited, searched, or otherwise
5 manipulated by the user by inputting instructions on keyboard 34. Other types of user input devices can also be used in conjunction with the present invention. For example, pointing devices such as a computer mouse, a track ball, a stylus, and/or a tablet can be used to
10 manipulate a pointer, such as a cursor, on a screen of a general-purpose computer.

Computer system 10 can also be used as a host or source computer for creating and/or providing the optimized documents of the present invention to "client"
15 (receiving) computer systems that download the documents. Alternatively, the host computer can be a file server or other type of mass storage apparatus.

Figure 2a is a diagrammatic illustration of a display screen 22 showing displayed visual
20 representations from a portable electronic document. The present invention is primarily directed to creating and downloading pages of portable electronic documents. A "portable electronic document" is a collection of data which includes objects which have been stored in a
25 portable electronic document language. The document is organized and stored in a "document file", which can be a storage unit such as a file, data structure, or the like. Portable electronic documents can be stored in a variety of different languages and formats. Herein, the portable
30 electronic document is described with reference to the Portable Document Format (PDF) by Adobe Systems, Inc. of Mountain View, California, or similar types of formats. PDF is a "page-based" format, in that a document includes a number of pages and is typically presented to a user on
35 a page-by-page basis, i.e., the user typically views one

- 21 -

page (or a portion of a page) at a time on a display screen. Other page-based document formats with similar document structures can also be adapted for use with the present invention.

5 Portable electronic document languages, such as PDF, typically store data as objects. An "object", as used herein, is a logical software unit comprising data and processes which give it capabilities and attributes. For example, an object can be queried as to its type and
10 can return such data as the number of words that it contains, its location in coordinates (e.g., location of the object's bounding box (BBOX)), etc. Objects can contain or refer to other objects of the same or of a different type. Objects can also be used to project
15 images on a screen according to their object type. There are many well-known texts which describe object oriented programming. Examples of object types used in typical PDF files include page objects, page contents objects (including text characters, words, etc. and/or graphical
20 objects such as polygon shapes, commands, etc.), image objects (e.g., bitmaps), font objects, and user-specific objects. Some objects can include direct references to other objects, and/or information used to display the object. In a PDF document, for example, a page object
25 can reference a page contents object that includes commands and text characters, where the text characters are provided as character codes representing the identity of the text characters. The page contents object can include the location to display the text, such as
30 coordinates used to display a bounding box around the text, and other information, such as the size and orientations of the characters. The page contents object can also refer to "shared objects", such as fonts (described below). Portable electronic documents, as
35 defined herein, differ from normal ASCII text files,

- 22 -

since ASCII text documents or files include only ASCII codes of characters and no other display information. Herein, objects may be "on" the page described by a page object, which is equivalent to an object being "referred to" by the page object.

In Figure 2a, display screen 22 shows a displayed page representation 40 derived from data received from a portable electronic document. Page representation 40 is displayed on the screen in a view window 39 by a viewer program from objects stored in the portable electronic document, and are typically derived from rendered bitmaps as is well known to those skilled in the art. A viewer application program ("viewer") running on computer system 10 can typically display an entire page, or a portion of a page, of a portable electronic document. The data for the page is requested by the viewer and received from a storage device or other computer. The data for page 40 may be downloaded from a host computer and displayed. The viewer may also offer features such as menu headings 41, selection buttons 43, and a table of contents or "bookmark" view 45. These features allow a user to manipulate the received data and view the page data according to user preferences. For example, menu headings allow a user to view, copy, load, save, search, or similarly manipulate the downloaded page. Selection buttons 43 similarly allow a user to view or manipulate the document in different ways by zooming, selecting the next page of the portable electronic document, etc. Bookmark 45 allows a user to select and display a particular portion of the document that the user (or a different user) has specifically marked and labeled with text (or, alternatively, graphics). For example, different chapter headings can be displayed as labels in bookmark 45 so that when the user selects a chapter, the first page of that chapter is displayed in view window

- 23 -

39. A suitable viewer application program for use with the present invention is Acrobat[™] Reader or Acrobat[™] Exchange, available from Adobe Systems, which can read, write, or manipulate page data stored in PDF format from a document file and display the data in a page-based format. Other viewers and file formats can be used in other embodiments.

Page representation 40 can include several different types of visual representations, including text 42, graphics 44, images, and links 46. Text 42 is derived from character codes and font objects stored in the document file. Text 42 can be rendered into a bitmap for display on screen 22, as is well known to those skilled in the art. Graphics 44 can also be rendered from coded shape primitives, such as lines and rectangles, and displayed. Images (not shown) are typically bitmap images, such as a scanned or digitized picture, and can be displayed on screen 22 by methods well known to those skilled in the art. Links 46 portray a topic or idea that can be accessed by the user and, for example, can enclose special text, graphics, or images to distinguish them from normal objects. Links 46 may be selected by the user to display a different portion of the portable electronic document that is related to the topic or idea portrayed by the link. Also, links 46 can be linked to other electronic documents that include the topic represented by the link to provide access to those other documents. For example, on many existing network services, links to many different documents available on the network are included in electronic documents. Page representation 40 can also be displayed on a sheet of paper output by printer 24. The computer determines the font, size, color, or other appropriate information for each object to be displayed by examining the associated font objects, color maps, size, and other information in

- 24 -

the PDF document.

Figure 2b shows a diagrammatic illustration of a display screen similar to the screen shown in Figure 2a. In Figure 2b, the bookmark view 45 has been replaced, due to a preference of the user, by a "thumbnails" window 48. Thumbnails window 48 displays page icons 50 (or "thumbnails"), each of which represents a separate page in the viewed portable electronic document. Icons 50 allow random access to any of the pages of the document, i.e., a user may select an icon 50 to display the corresponding page 40 in view window 39. The currently-displayed page may have a highlighted label, such as icon 51.

Figure 3a is a diagrammatic illustration of a non-optimized document file 54. In the described embodiment, a document file having data stored in the PDF language is referenced as the main example. In other embodiments, the document file can have a different page-based format. A portable electronic document is typically stored as a non-optimized document file 54 as shown in Figure 3a when written to a file or memory using normal processes of the prior art. A non-optimized document file includes page contents data 56 that is usually stored in a disjointed manner within the file 54. The page contents data includes data for text 42 and graphics 44 for a single page. In addition, the page contents data also includes other related data not shared by other pages, such as referenced (unshared) fonts, images, procsets, etc. For example, a particular page P has text/graphics page contents 56a placed at the beginning of the file 54, unshared image page contents 56b placed at a later position in the file discontinuous from contents 56a, a resource dictionary contents 56c discontinuous from contents 56b, and unshared font object contents 56d placed near the end of the file. All of

- 25 -

these contents need to be read by the viewer to display the entire page P. Other pages of the document have page contents data 56 stored in file 54 in a similarly disjointed manner.

- 5 In addition, document file 54 may include a cross reference table 58, which provides a table of each object in the document file 54 indexed to the location (offset) in the file where the object is located. The Portable Document Format Reference Manual, Adobe Systems Incorporated, Addison-Wesley Publishing Company, New York, 1993, describes the cross reference table, pages tree, and other features of the PDF file format, and is hereby incorporated by reference herein. The cross reference table 58 can be placed anywhere in the file, e.g., at the end of the file; or, parts of the cross reference table can be located in different portions of the file. Finally, the non-optimized document file 54 typically includes shared objects 60 and special objects 61 (described below), which are typically located throughout file 54 in a disjointed manner. Shared objects may be referenced by multiple page contents objects in the file, and can include font objects, color maps (or "color spaces"), and other objects which are necessarily referenced to influence the appearance of an object when displayed. Shared objects may also include any objects appearing on multiple pages and user-defined shared objects. The "shared objects" referred to herein may not be actually be shared in a particular instance; these objects, however, can be potentially shared. For example, a font might only be used by one page and may not be shared by other pages in a particular document, but it can potentially be shared by other pages. As is well known in the art, a PDF file typically relates objects in a "page tree" structure, where an object may refer to a child or descendant object. For example, a

- 26 -

page object may refer to page contents 56 (child) objects, while the page contents objects further refer to (child) shared objects 60.

In a process of displaying a page of document file 5 54, the document file 54 may be located on a separate host computer. When a user requests that the viewer shown in Figures 2a and 2b display a particular page P of the document file on display screen 22, the viewer first establishes a connection to the host computer to access 10 the desired document file. If the page contents are organized as shown in Figure 3a, the viewer first downloads a designated amount of page contents 56a. After several possible connections to download contents 56a, another connection must typically be made by the 15 viewer to access contents 56b, and so on. In addition, if a page P's contents reference an object such as a font, that font may have to be downloaded before the page contents which require that font can be displayed. This all contributes to a long delay for the user before any 20 portion of a page is displayed by the viewer.

Figure 3b is a diagrammatic illustration of an optimized document file 62 of the present invention. The data in file 62 has been organized to minimize the amount of time to download a page from a host computer and 25 display the page by the viewer as shown in Figures 2a and 2b. Document file 62 includes a range table 66 stored at the beginning of the file, and a first page portion 64 of the cross reference table also stored near the beginning of the file. The page contents 56 are grouped and stored 30 contiguously, so that a contiguous amount of the first page's contents 56 is stored, followed by a contiguous amount of page 2's contents 56, and so on until the end of page N's contiguous contents 56, where N is the last page number in the document. All shared objects 60 may 35 be stored after the page contents data. Special objects

- 27 -

61 that are not required for displaying a page may be stored after shared objects 60. In the described embodiment, special objects include page icon (thumbnail) objects, bookmark objects, page tree objects, and the like. The cross reference table 58 may be stored after the special objects, and a page offset table 68 of the present invention may be stored at the end of the file. The page offset table provides the locations of pages in the document file 62 and shared object information to the viewer, as described subsequently. In alternate embodiments, the page offset table can be stored in other areas of the file. One such alternative embodiment is illustrated in, and discussed in reference to, Figure 14. The function of the organization of data as shown in Figure 3b is described subsequently. An application program such as the viewer shown in Figures 2a and 2b may include an option to save a viewed document (or a document on a specified storage device) as an optimized document file as shown in Figure 3b. If it does, the user could select whether to save a document in optimized or non-optimized format.

In alternate embodiments, the document data can be written in different locations of optimized document file 62. For example, the range table 66, cross reference table 58 or 64, or page offset table 64 can be placed at particular locations in the file 62, and the viewer can read those specific locations when particular data needs to be downloaded.

CREATING AN OPTIMIZED DOCUMENT FILE

Figure 4 is a flow diagram illustrating a process of the present invention for creating an optimized document file 62 of the present invention from a non-optimized document. The "non-optimized document file" can be stored as a file on a storage device, or can

- 28 -

be partially or wholly stored in memory (such as RAM 16) of a computer system 10, such as during a viewing process as shown in Figures 2a and 2b. The non-optimized document, in the described embodiment, is in "normal" PDF format. The process 70 can be initiated by a user who wishes to save the non-optimized document as an optimized document file of the present invention. For example, a non-optimized document can be loaded into the RAM of a computer system and then saved to hard disk or other computer-readable storage medium as optimized document file 62. For example, PDFWriter or Acrobat Distiller from Adobe Systems are used to write PDF files from application programs, and can be used to implement the process 70. The saved optimized file 62 can, for example, be made available on a host computer or server ("host" or "host computer") to client computer systems that may request the document for downloading. The process of the present invention for downloading optimized document file 62 is described in greater detail with respect to Figure 10.

The process begins at 72. In step 74, an internal list of objects and lists of shared objects are created from the non-optimized document file. These lists help the process to organize the objects of the document for grouping the objects in the more optimized configuration of the present invention. The lists of shared objects are used to place shared objects in the document after the page contents objects. Step 74 is described in greater detail with respect to Figure 5. In next step 76, the document information, including page content information 56, shared objects 60, and special objects 61 are written to the optimized document file according to the internal list organized in step 74. The process essentially retrieves each successive object name or "ID" in the internal list and writes the object to the

- 29 -

optimized file in the same order. This has the effect of organizing the page contents 56, shared objects 60, and special objects 61 for all the pages of the document as shown in Figure 3b. In addition, space may be allocated
5 at the beginning of the optimized file to store the cross reference table portion 64 and range table 66, which are described below. Also, information for the cross reference table 58 describing the locations of objects in the file is stored in memory as the objects are written
10 in this step.

In an alternate embodiment, the page contents and shared objects can be stored in optimized document file 62 in an interleaved order, where portions of page contents are followed by shared objects referenced by
15 those portions. In one embodiment, the data is read or downloaded in an interleaved order, but is not stored in the interleaved order. Interleaving is described in greater detail with reference to Figures 9 and 12. In an alternative embodiment, only the first page's contents
20 are physically stored in an interleaved order in the optimized file so that range table 66 need not be downloaded, which may save a connection to (or a transaction with) the host computer and reduce initial downloading time of the file.

25 In next step 78, the cross reference table 58 (or equivalent structure using other file formats) is written to the file, and in one embodiment, to the end of the file after the special objects 61, as shown in Figure 3b. The cross reference table 58 is a listing of objects in
30 the document and the offsets (e.g., in bytes) from the beginning of the file for the start of each object, and allows random access to the objects in the document (the cross reference table, however, does not provide the types of objects or pages where objects are located). A
35 "trailer" for a PDF file can also be written is this

- 30 -

step. The information for the cross reference table is determined as each object is written to the optimized file in step 76, as explained above. The formation of cross reference tables in PDF files are well known to those skilled in the art.

In step 80, the page one portion 64 of cross reference table 58 is written to the optimized document file 62, and in one embodiment it is written near the beginning of the file (leaving room to store range table 66). This portion 64 of the cross reference table 58 refers to objects on the first page of the electronic document and is placed near the beginning of the file so that page one information can be retrieved as soon as possible in a downloading process. The contents of the first page can thus be displayed immediately upon receiving those contents when downloading document file 62 (as in Fig. 10).

In step 82, the page offset table 68 of the present invention is created and placed in the optimized file. In one embodiment, it is placed near the end of the optimized file 62 after the cross reference table 58. The information in the page offset table allows any page of the electronic document to be quickly accessed and downloaded. The page offset table may be included in the optimized file 62 so that a separate offset table file does not have to be downloaded, processed, updated, or associated with a page contents file. However, in alternative embodiments, the page offset table and, optionally, other index tables may be stored in one or multiple secondary files, residing optionally on separate host computers. In one embodiment, a pointer to the page offset table is included in the range table 66 (described below) near the beginning of the file 62, so that the page offset table can be accessed after the first page one information has been downloaded in a

- 31 -

downloading process. Since the page offset table is not needed to display page one of the document, it is not placed at the beginning of the file. Alternatively, the page offset table can be provided at other places in the optimized file 62, including at the beginning. The page offset table is described in greater detail with respect to Figure 8.

In step 84, the range table 66 (and, optionally, a header) is written at the very beginning of the optimized file 62. The range table provides the offsets and lengths for the page contents and shared objects of the first page of the document, and thus functions somewhat like the page offset table 68. The process of writing the range table is described in greater detail with respect to Figure 9. Other necessary information can also be written in this step; for example, PDF files store a pointer to the contents of the document file at the end of the file after the "trailer." The process is then complete at 86.

Figure 5 is a flow diagram illustrating step 74 of Figure 4, in which an internal list of objects and lists of shared objects are created from the non-optimized portable electronic document. The process begins at 88. In step 89, any inheritance data from parent objects are copied into children objects down the page tree, as appropriate. This step is implemented if, as in PDF files, some children objects do not include certain needed data and instead refer to and "inherit" this needed data from a parent object, e.g., display commands, orientation commands, etc. Since the present invention records object data and may not be able to reference a parent object easily, any inheritance data from parent objects is copied into children objects that need such data in step 89.

- 32 -

In step 91, the page tree objects are added to a separate page tree internal list, except for page tree objects for the first page. Only the first page's page tree objects need to be initially downloaded in a

5 downloading process, so the other page tree objects can be organized in a separated list in step 91 and, later, written near the end of the file (described below). In step 92, a variable P is initialized to 1 and checked if it is less than or equal to the number of pages in the

10 document. If so, in step 94, the page object for page P is retrieved from the non-optimized document file (or non-optimized document stored in memory) and is written to an internal list. A page object, as typically defined in a PDF file or similar format, is an object that refers

15 to other objects which are included and displayed collectively on the page. Thus, by retrieving the page object for page P, the process also indirectly retrieves references (pointers) to page P's page contents object and any other referenced objects. Step 94 finds the page

20 object for page P in the non-optimized file by use of the cross reference table from the non-optimized file (which can be utilized by the viewer). The page P object identification (ID) is written to an internal list, stored in RAM 16 or other storage. In PDF and other

25 types of files, an object typically includes an object ID, which is a number or other identifier that uniquely identifies the object within the document file. A number identifier for the page P object is written to the internal list in this described embodiment.

30 In step 96, if page icons 50 ("thumbnails") are implemented in the viewer embodiment as shown in Figure 2b, the page icon object for the current page P is typically added to a separate page icon list if the initial displayed view does not include page icons. In

35 viewers such as Acrobat, however, the user can save a

- 33 -

file designating the document to be displayed with a page icon view as shown in Figure 2b when initially downloaded and displayed. If the current non-optimized document has been so designated, the page icon objects are needed to display the first page. Therefore, in one embodiment, step 96 can be skipped and the page icon objects can be added to the internal list after all first page objects have been added (i.e., a negative determination of step 108 for the first page (e.g., $P = 1$)). Alternatively, a page icon for a specific page can be stored after its page contents. Other special objects which are not necessary for displaying the document can also be added to specialized internal lists in step 96 so that these objects can be written to the end of the document file.

Herein, "special objects" refer to page icon objects, bookmark objects, page tree objects, and any similar types of objects that are usually not necessary to display a page and can be provided near the end of the optimized document file.

In next step 98, the next object "on" page P is retrieved in a designated order, i.e., the next object referred to by the page P object in the designated order of objects. The "designated order" of objects is the order of objects in which the provider of the optimized file (or the implementor of process 74) desires to be downloaded and displayed when accessing a page from the file. Thus, when downloading the optimized file 62, certain types of objects can be displayed first while other types of objects are still being downloaded. For example, it is typically desirable to display the text (i.e., page contents object) on a page before the images on the page when the page is downloaded, since the user who is downloading the page can read the text while the images are still being downloaded. From the text content, the user can quickly determine if it is worth

- 34 -

his or her time to wait for an image to be downloaded and be displayed, or if that page download should be interrupted and a new page downloaded.

In one embodiment, the designated order of objects is links, non-image "resource objects" (i.e., non-image objects in a resource dictionary), page contents objects, image objects, bookmark objects, user-defined objects, and other objects. This order allows links to be downloaded and active first. Since links are typically implemented as rectangular (or other shaped) "bounding boxes" which enclose text, graphics, images, or other objects, the links are advantageously first in the designated order so that when an enclosed object of a link is later displayed, the link will already be receptive to user inputs. Other types of objects that may be added to a document by a viewer are ordered with links in the designated order. Non-image "resource objects" are ordered next, which include shared objects such as font objects, color map objects, and the like. Font objects provide the data to determine how text will appear, and color map objects map colors to different display output devices, as is well known to those skilled in the art. The font and color map objects are needed to display text and graphics, and thus should be downloaded close in time to the page contents objects which refer to them. In PDF files, these types of objects are typically located in a "resource dictionary" which is located in the page contents object 56 of a document file (or as a separate object) and is used for decoding page contents to map objects with object references. Image objects can also be referenced in the resource dictionary in typical PDF files; however, they are ordered later in the designated order. The page contents objects (text and graphics) are ordered next in the designated order, which allows a user to quickly download and view the

- 35 -

text/graphics and determine the subject content of the page. The image objects are ordered next, near the end of the designated order, since they are usually the largest size objects and require the greatest downloading
5 time. Objects referenced by a bookmark object, as shown in Figure 2a, are next in the order. These objects should be downloaded so that they can be accessed by the bookmark, and are a special case, as explained below. User-defined objects and any other types of objects are
10 ordered last.

Other designated orders of objects can be implemented in other embodiments. For example, links can be downloaded after text and graphics objects. Alternatively, the user who is creating the optimized
15 file 62 can be offered an option of inputting a desired designated order of objects.

The next object in the designated order of objects is thus retrieved in step 98. That is, if this is the first time step 98 is implemented, a link object is
20 retrieved. Once all the link objects have been retrieved and added to lists in step 102 of the current process, an object next in the designated order (e.g., resource objects) is then retrieved in step 98, and so on.

In step 100, the process determines if the
25 retrieved object has already been examined for a different page or if the retrieved object is a forced shared object. If the retrieved object has already been examined for different page, then this object is designated a shared object, i.e., the object was referred
30 to by an earlier examined page/page contents object or the current page/page contents object and was already retrieved and processed. Any object can potentially be a shared object; for example, a page contents object or an image object in the document file can be referred to by
35 two or more different pages. The process can check if an

- 36 -

object is shared by examining the internal list that has been created so far in the current process. If an object ID is found that is identical to the current object's ID, then the object is considered to be a shared object.

5 Also, in one embodiment, in step 100, if the retrieved object is a forced shared object, then the object is automatically forced (designated) to be a shared object, regardless of whether the object is actually referenced by multiple pages or not. Herein,
10 "forced" shared objects include such resource objects as font objects and color map objects, but do not include resource objects such as image objects or procset objects. Font and color map objects are forced to be shared objects because they are required in the decoding
15 of page contents and may, in particular embodiments, be advantageously interleaved in the downloading process, described below. "Procsets" are used for printing purposes, as is well known to those skilled in the art, and are need not be automatically forced to be shared
20 objects (images and procsets can be shared objects if they are referenced by multiple pages).

 In addition, the user can provide his or her own non-standard shared objects that can be referenced on multiple pages. For example, a user could provide a
25 dictionary table object in the document that is referenced by multiple pages. Such user-defined shared objects are not necessarily referred to by and are not required to display page contents such as text, and are therefore referred to herein as "non-contents shared
30 objects."

 If the retrieved object is not a shared object, then step 102 is performed, in which the object ID of the retrieved object is added to the end of the internal list. The internal list thus has an order of objects
35 including a page object followed by all the objects (in

- 37 -

the designated order) referenced by that page.

In step 103, special objects are processed if predetermined conditions are met. For example, in the described embodiment, bookmark 45 objects are a special
5 type of object that are processed in step 103 if the current page is the first page ($P = 1$), and if image objects on the current page have all been processed (i.e., bookmark objects are after image objects in the designated order). If these conditions are met, then the
10 process also checks if the bookmark view 45 is to be displayed when the document is initially opened, i.e. if the bookmark view of Figure 2a is the default initial view. If so, the ID's of a predetermined number of bookmark objects (e.g., 60) are added to the internal
15 list. This number is the number of bookmark objects that would be initially shown in the bookmark view 45 as shown in Figure 2a. The remainder of bookmark objects are then added to a separate bookmark internal list, similar to the separate page icon list described above. If the
20 bookmark view of Figure 2a is not the default initial view, then all of the bookmark objects are added to the separate bookmark list in step 103. This step allows objects displayed in the bookmark view to be grouped with the first page's contents data so the bookmark objects
25 can be downloaded and displayed quickly with first page data in a downloading process (if the initial document display includes the bookmark view). The process then continues to step 108, detailed below.

If the retrieved object is a shared object in step
30 100, then step 104 is implemented, in which the shared object is processed and a sharing pages list is created. This step is described in greater detail with respect to Figure 6. In next step 106, the object ID of the shared object is added to the shared object list, which is
35 similar to the internal list except that it includes only

- 38 -

shared object ID's of objects that have been examined by process 74.

Figure 5a is a diagrammatic illustration of shared object list 148 as processed in step 106 of Figure 5.

5 The shared object list 148 includes a node 150 for each shared object found in the process of Figure 5. Each node 150 includes an object ID field 152 for storing an object's object ID, and a contents flag field 154 for storing a contents flag, as detailed subsequently.

10 After step 102 or step 106, step 108 is implemented, where the process checks if there are any additional objects on page P that have not been examined. If so, the process returns to step 98 to retrieve the next object in the designated order of objects. If there
15 are no additional objects, step 110 is implemented, in which the number of objects on page P is stored in memory. This number can be determined by counting all the newly-added object ID's in the internal list. The process then returns to step 92 to increment the variable
20 P and again check if P is less than or equal to the number of pages in the document.

Once all of the pages of the non-optimized document have been processed by the aforementioned steps, the process continues to step 112, in which a sharing
25 pages list is completed and all sharing objects and other objects are incorporated into the internal list. The sharing pages list is created in the processing of shared objects of step 104. After the shared and other objects have been incorporated into the internal list, then the
30 internal list is ready to be used to write all of the objects in the optimized order in step 76 of Figure 4. Step 112 is described in greater detail with respect to Figure 7. The process is then complete at 114.

Figure 6 is a flow diagram illustrating step 104
35 of Figure 5, wherein the retrieved shared object is

- 39 -

processed. The process 104 begins at 116, and, in step 118, the process checks if the retrieved object is the first shared object on page P. This is determined by checking if page P is already in the sharing pages list.

5 If page P is already in the sharing pages list, then a shared object was already found to be referenced by page P, and step 122 is implemented. If page P is not in the sharing pages list, then, in step 120, page P is added to the sharing pages list.

10 Figure 6a is a diagrammatic illustration of a sharing pages list 136 of the described embodiment. This list is created and referred to during the process 104 of Figure 6. As described, the list is provided as a linked list, the implementation of which is well known to those
15 skilled in the art. Other implementations of the list can also be provided. A page is represented by a page node 138, where each page node is linked to another page node, in sorted numeric order of the pages, for example. If a page node exists in list 136, then that page
20 includes shared objects, as determined by steps 118 and 120 of Figure 6.

Referring back to Figure 6, if the object is not the first shared object referenced by page P, or after step 120, then step 122 is implemented. In step 122, the
25 process checks if the retrieved object is in the object list of page P. As shown in Figure 6a, an object list 140 is referenced by a page node 138 if that page references shared objects. An object list 140 includes a number of object nodes 142, each of which represents a
30 shared object referenced by the page of the corresponding page node. In step 122, the process checks if the retrieved shared object is already in the object list 140 of page P. If such a condition exists, it indicates that multiple references to the shared object are present on
35 page P. Since only the first instance of the shared

- 40 -

object on a page is needed for the present invention, step 126 is then implemented, as detailed below. If, however, the retrieved object is not in the object list 140 of page P, then, in step 124, an object node 142 is
5 added to the object list 140 and a sharing ID is added to an ID field 144 of the object node (object node 142 also includes a fraction field 146, detailed below). The "sharing ID" is an identifier for the object which uniquely identifies the object in the sharing pages list.
10 The sharing ID is may be a number, n, that indicates a shared object is the nth shared object that has been found in the document. For example, a sharing ID of "0" indicates that an object is the first shared object found in the document, a "1" is the second shared object found,
15 etc. The sharing ID is not the same as the object ID, since a sequential object ID, m, would indicate that the object is the mth (shared or non-shared) object in the document. Object node 142 and the sharing ID would naturally be added to the object list 140 in a sorted
20 numerical order according to sharing ID's.

Step 126 is then implemented, in which the process checks if the retrieved object is the first shared object on the original page, i.e., if the original page is already in the sharing pages list 136. The "original
25 page" is any other previous page that also references the shared object. There may not be an original page if, for example, the retrieved shared object is a forced shared object (e.g., a font). The process determines the original page by, for example, checking a table which
30 logs each object and the page that each object is located on. If the retrieved object is not the first shared object on the original page, then step 130 is performed, detailed below. If the retrieved object is the first shared object on the original page, then in step 128, a
35 page node 138 corresponding to the original page is added

- 41 -

to the sharing pages list. New nodes may be added to list 136 in a numerical sorted order by page number. For example, a shared object designated by node 142a is found on page 9, and is assigned a sharing ID of "1". Page 1
5 also includes the same object, so that node 142b is added to the object list 140 of the page node for page 1. The process then continues to step 130.

After a negative determination of step 126, or after step 128, step 130 is performed, in which the
10 process checks if the retrieved object is in the object list(s) 140 of the original pages(s). This step is substantially similar to step 122, above. If the object is already in the original page's list 140, then the process continues to step 133, described below. If the
15 object is not in the original page's list 140, then, in step 132, an object node and sharing ID for the retrieved object are added to the object list 140 of the original page. The process then continues to step 133.

In step 133, the process checks if the retrieved
20 object is a forced shared object or if the object is not in the object list of page P. If either condition is true, step 131 is performed, in which steps 104 and 106 of Figure 5 are recursively performed for all of the children objects referenced by the retrieved object (if
25 any), including children objects of other children, etc. This step forces children objects of parent shared objects to also be shared. Such children objects can include, for example, widths of characters for a custom font, etc. After step 131, or if neither of the
30 conditions of step 133 are true, then the process is complete at step 134.

Steps 118-124 can be implemented as a function which examines an object passed to the function through a function call. This same function can then perform steps
35 126, 128, 130, and 132 by passing the original page to

- 42 -

the function through a second call. Figure 7 is a flow diagram illustrating step 112 of Figure 5, in which the sharing pages list is completed and shared and other objects are added to the internal list. The process 112
5 begins at 156. In step 158, the next page node 138 in the sharing pages list 136 is retrieved. If this is the first time performing step 158, then the first page node in the list 136 is retrieved. In next step 160, the process determines if there are any more page nodes in
10 the list 136 to examine, e.g., if a null or end of list symbol was retrieved in step 158 to indicate all page nodes have been examined. If there are no more page nodes to examine, step 170 is performed, detailed below. If a new, unexamined page node was retrieved in step 158,
15 then, in step 162, the names of the resource objects which are also shared objects for the selected page are retrieved from the resource dictionary utilizing shared object list 148 (a resource dictionary is typically associated with each page object). Only certain types of
20 resource objects are retrieved which are desired to be interleaved in the downloading process. In the described embodiment, these desired resource objects include font objects, color map objects, and (shared) image objects. Each resource object typically has a "name", which is an
25 identifier for the object so that it may be referenced to the actual object data. For example, font objects can have a name such as "f1" or "f2" which references a particular font object for the font of "Helvetica", "Times", etc.
30 In next step 164, the process searches for the retrieved shared resource object names in the page contents portions of the selected page. For example, if page 9 is the selected page, the process accesses the page contents of page 9 and searches for names (such as
35 "f1") from the resource dictionary found in step 162. In

- 43 -

step 166, for each found resource object, the fraction number of the page contents that includes the found resource object is written into the object list, if this is the first found occurrence of the resource object on the page. For example, fractions can be designated a size of eighths, such that a fraction is $1/8$ the size of the page contents, and the fraction number for the first fraction is 0 (zero), representing $0/8$ to $1/8$ of the page contents data of the page. Larger or smaller fractions can be designated in other embodiments. The fraction size is used when interleaving and can depend on the desired amount of page contents data to be downloaded before the shared objects referred to by that contents data are downloaded, as explained subsequently. When a resource object's name is found in the page contents, the process has found a reference or pointer to the resource object. For example, text in the page contents may include an identifier referring to a font object. The particular fraction of the page contents in which the resource object was named is then written into the object list 140 of the current page node. This fraction number is written into fraction field 146 of an object node 142 and is the dividend of the fraction, with the divisor being a predetermined number, which may be found in the document file header. For example, for object node 142c of list 136 in Figure 6a, a fraction number (dividend) of "7" indicates that this shared object occurred in the $7/8$ to $8/8$ portion of the page contents data of page 22, i.e., if the page contents were divided into eighths, the object would occur in the last eighth, where "8" is the predetermined divisor. A fraction number is similarly written for each resource name searched and found on the page. Also, the fraction numbers written into parent shared objects are also written into any children objects of those parent shared objects (children objects are

- 44 -

described in step 131 of Figure 6) so those children can be interleaved with the parent objects. In one embodiment, image resource objects are automatically assigned a high fraction number so that images will be
5 downloaded late in an interleaved downloading process, as described below. In an embodiment using a method of progressive rendering of fonts, described in reference to Figure 13a, font descriptor objects are given correct
10 fraction numbers, so that they will be downloaded close to the font reference in the interleaved downloading process, while font data objects are given a high fraction number, so that they will be downloaded toward the end of the page.

Also in step 166, the nodes of each object list
15 140 may be reordered so that the shared objects are provided in an order from first to last occurrence on the page. In an embodiment where image objects are ordered from first to last occurrence separately from other shared objects on the selected page, these may be added
20 to the end of the object list for the page.

In step 168, the contents flag associated with each found resource object is set in the shared object list 148. As shown in Figure 5a, the flag field 154 holds the contents flag for each shared object. This
25 flag is set to "1" if the shared object was found in step 166. The contents flag indicates which shared objects should be interleaved with page contents when the optimized file is downloaded, as explained subsequently. Shared objects that do not have the contents flag set
30 will not be interleaved in the downloading process of Figure 10. (Non-contents objects, such as user-defined objects, procsets, and resource dictionaries, even if shared, do not have the contents flag set.)

After step 168, the process returns to step 158 to
35 retrieve another page node 138 from sharing pages list

- 45 -

136. Once all the page nodes have been examined, the process continues from step 160 to step 170, where the shared object ID's from the shared object list 148 are appended to the end of the internal list. In step 171, 5 the bookmark objects on the bookmark internal list, page icons on the page icon internal list, page tree objects on the page tree internal list, and any other required objects are added to the end of the main internal list. In next step 172, any duplicate shared object ID's in the 10 front portion (i.e., portion before the shared objects portion) of the internal list are removed from the internal list. The shared object ID's in the shared object list 148 are compared with the object ID's in the front portion of the internal list, and any matches from 15 the internal list are removed. In addition, in step 172, the total number of objects for a page is decreased by the amount of matched objects so removed. In step 174, the shared object ID's may be reordered, using the contents flags set in step 168, so that shared objects 20 referenced by page contents are ordered first. This order allows the contents shared objects to be grouped and the page offset table to require less storage space than if the shared objects were not reordered. The process is then complete at 176.

25 Figure 8 is a flow diagram illustrating step 82 of Figure 4, in which the page offset table of the present invention is developed and stored in the optimized document file. The process of Figure 8 generates one example of a page offset table, which can have other 30 formats and forms in alternate embodiments. The process begins at 180, and, in step 182, the process determines the length of each page, in bytes or similar measures of content. This can be determined by examining the number of objects on a page and determining the byte offset 35 amount between the first object on the page and the first

- 46 -

object of the next page. In step 184, housekeeping information is written into the page offset table 68. An example of a page offset table 68 as referred to by process 82 is shown in Figure 8a. The housekeeping information can include the number of pages in the electronic document, the least number of objects on a single page, the location of the cross reference table (e.g., the offset in bytes from the beginning of the file), and the number of shared objects in the document.

10 This housekeeping information is used to create page information to help download portions of the document, as detailed in the downloading process of Figure 10.

If there are shared objects in the file, then shared object housekeeping information is written to the page offset table in step 184. The shared object housekeeping information includes the number of non-contents shared objects in the document, the least size of a shared object, and the size of the dividend for the fraction size (e.g., 3 bits).

20 In step 188, the number of objects on each page as determined in step 110 of Figure 5 and step 170 of Figure 7 and the page length information as determined in step 180 are compressed and written to the page offset table, shown as information 189 in Figure 8a. In step 190, for each shared object in the document, the length of the shared object (e.g., in bytes, determined by offset comparison) and the signature for the shared object (if applicable) are stored in the page offset table. A signature is an identifier that allows caching of shared objects to be used safely in the downloading process.

30 Signatures are used to avoid mistaking resources (such as fonts) that may be different but nevertheless may go by the same name. A signature may be calculated from the resource itself by any method with a sufficiently high likelihood of giving different values for different

35

- 47 -

resources. Thus, a shared object can be stored in a disk cache after it is first downloaded for use whenever that shared object is accessed or referenced by page contents in the same or even in a different document file.

5 Next, in step 192, for each page P of the document, additional information is determined and collected for compression for page P, including a shared object flag, the number of shared objects on page P, sharing ID's, approximate page contents fraction
10 information, and the divisor value used for the fraction size. The shared object flag indicates that a page P references shared objects. Shared object ID's are stored for the shared objects being used by that page. The approximate contents fraction information is estimated
15 from the sizes of page contents objects as a fraction of page length; e.g., this approximate fraction can be designated as "1/8". The divisor value is, for example, "8" if the fraction size is determined to be eighths. After the information in step 192 has been determined for
20 each page of the document, the process continues to step 194, where the information for all the pages may be compressed by well-known techniques and written to the page offset table in the optimized document file 62. The process is then complete at 196.

25 The page offset table 68 is intentionally made small and compact in size, where data is compressed when possible. Since the page offset table is additional data that is not normally downloaded in non-optimized documents, it is desirable that the additional data be a
30 small as possible so that the user does not have to wait any extra length of time. In addition, since the page offset table is compact, it is less "noticeable" (through time delays) in a downloading process than other much larger structures, such as the "page tree" normally
35 included in a PDF document.

- 48 -

In alternate embodiments, the page offset table can include different or additional data, or the data can be stored in different formats. For example, the sharing pages list 136, internal list, and other lists generated
5 can be stored directly (and inefficiently) in the page offset table in some embodiments.

Figure 9 is a flow diagram illustrating step 84 of Figure 4, in which a range table 66 for page one of the electronic document is written at the beginning of the
10 optimized document file. The range table provides information so that the page one data of the electronic document may be located in the optimized file and quickly downloaded and displayed. It is assumed that page one is the "first page", i.e., desired to be displayed first, as
15 a default, when beginning to download a page-based electronic document. In other embodiments, a different page can be the default first page that is displayed.

The process begins at 200. In step 202, housekeeping information is written to the range table.
20 This housekeeping information is similar to the housekeeping information described with reference to the page offset table above. Figure 9a shows a diagrammatic illustration of a range table of the present invention. The housekeeping information includes the number of
25 ranges of data that are downloaded for the first page. These ranges can include interleaved portions of page content and shared objects (described below). The shared pages list 136 can be referenced to determine how shared objects are to be interleaved with page content by
30 examining the fraction numbers in the object nodes 142 and appropriately interleaving the shared objects with the page content. Alternatively, no interleaving is applied, and the first range is an entire page contents object 56, while any additional ranges are shared objects
35 referenced by the page contents. A version number for

- 49 -

the range table/page offset table in the current file may also be stored as housekeeping information. Also, a signature flag for each range can indicate if the range (object) has a signature; the signatures can also be
5 listed in the range table. Signatures can be used for caching shared objects, as described above.

In step 204, a pointer to the page offset table 68 is written into the range table. This pointer allows the page offset table to be located in the optimized file and
10 downloaded after the range table and first page have been downloaded. In next step 206, the process checks if the total first page size is less than a predetermined minimum size. If the page is less than the minimum size, then no interleaving of page contents and shared objects
15 is desired, since the page contents are so small that no advantage in downloading speed may be gained by the interleaving; in fact, the downloading speed may be slower when the page is below the minimum size and interleaving is provided, creating a longer wait for the
20 user to view the page on the display screen. Interleaving of page contents and shared objects is described in greater detail with respect to Figure 12. For example, a minimum size of 4 kilobytes can be specified.

25 If the page is more than the minimum size, then, in step 208, the offsets and lengths of the page content are written into the range table interleaved with the offsets and lengths of the shared objects. That is, if a fraction of the page contents includes a reference to a
30 shared object, the referred shared object is ordered after that fraction of the page contents in the range table. The interleaved ranges determined in the housekeeping information in step 202 can be used. The process is then complete at 210.

- 50 -

If the page is less than the minimum size in step 206 (or if no interleaving is desired for a different reason), then, in step 212, the offset and length for the entire page contents are written to the range table, so that one offset and one range describe the page contents. In step 214, the offsets and lengths for the shared objects, and, in one embodiment, the offsets/lengths for the cross reference table entries needed for those shared objects, are written to the range table. They can be written, for example, in the order they have been stored in the object list 140. The process is then complete at 210.

Alternatively, other conditions can also be checked to adjust the arrangement of page content and shared object offsets in the range table. For example, if the page contents object is less than 3 kilobytes in size, then shared objects offsets can be placed after the page contents data with no interleaving.

DOWNLOADING AN OPTIMIZED DOCUMENT FILE

Figure 10 is a flow diagram illustrating a process 220 of downloading an optimized page-based document of the present invention from a host computer to a client computer system 10. It is assumed the user wishes to view the document in a viewer that can display a page or a portion of a page of the document. The "finder" described in the current process can be implemented within the viewer, or as a separate procedure or program instructions running simultaneously with the viewer on computer system 10.

The process begins at 222. In step 224, the process checks if the user has requested to view a document that is located, appropriate to the present invention, on a host (source) server or computer. If no request for a document is made, the process waits for

- 51 -

such a request (the computer system 10 or viewer can be performing other tasks for the user in the meantime, such as viewing local documents). When the user requests to view a downloaded document, the viewer connects to the
5 host computer in step 226 to download the header and the range table 66 for the optimized document file. The header includes information designating the file in a specific format, e.g., PDF, and the presence of the range table may be used to indicate the file is an optimized
10 file. The header and range table may be located at the beginning of the file so that they may be downloaded first. Other needed information is also downloaded at this time; for example, PDF viewers may require a pointer to the contents of the file that is located at the end of
15 the file. In the alternate embodiment in which the first page's contents are stored in an interleaved order in the optimized file, the range table need not be downloaded, and step 226 can be omitted.

In step 227, the process checks if the first page
20 of the documents has been requested to be viewed by the user in step 224. In one embodiment of the invention, when a user first requests to view a document, the first page will be automatically downloaded and displayed as a default. The user then requests a desired page after the
25 first page has been downloaded, as detailed below.

However, in other viewer embodiments, a user may be able to initially request to view a particular page of a new document before any part of the document is downloaded. In such an alternative embodiment, step 231 is
30 implemented, described subsequently. If the first page is always initially displayed, or the user requests the first page, step 228 is implemented.

In step 228, the first page data and page offset table are downloaded from the source file by the viewer
35 in another connection (or another transaction) using the

- 52 -

range table, and the first page is displayed. This first page data includes page contents and shared objects of the first page. The portion of the cross reference table for the first page is also downloaded at this time. The shared objects and page contents of the first page may have already been interleaved using offsets in the range table, as described with reference to Figure 9. Thus, in one embodiment, a portion of the page contents are downloaded, followed by any shared objects referenced by that portion (and cross-reference data for those shared objects), which allows that portion to be immediately displayed to the user. The user thus experiences very little waiting time to view at least some of the contents of first page. The page offset table 68 may be downloaded during this connection (or transaction) using the pointer in the range table 66 after the first page is downloaded. Thus, the page offset table may be downloaded early in the downloading process, immediately after (or, alternatively, before) the downloading of the first page, for example, so that other pages in the document can be randomly accessed and viewed. In some embodiments, the viewer may be selected not to display the first page when downloading a document.

In next step 230, the page offset table is examined to determine page information. The finder may be made responsible for examining and processing the page offset table, in which case the finder determines page information using information in the page offset table. The process of determining this page information is described in greater detail with reference to Figure 11. The process then continues to step 234.

In step 234, the process checks if the user requests to view a particular page of a document in the viewer. The user can request a page of the current document that was partially downloaded in steps 226, 228

- 53 -

and 230 (or step 226, 231 and 232) or a different document (described below). For example, to request a page of the current document, the user can select a page icon 51 as shown in Figure 2b to display a different page
5 of the current document. Or, the user can select a link to a different page in the current document or select a bookmark object listed in the bookmark view 45 as shown in Figure 2a. If no request to display a different page of a document is made, then the process continues to wait
10 for such a request at step 234 (other viewer or computer functions can be performed during step 234). If a request of the current document is made, the process continues to step 236, described below.

In step 234, the user may also be able to request
15 a page of a different document that has not yet been downloaded. For example, a link or other control in the viewer may be selected to provide access to a different document file available on the same or different host computer system. If the user requests a page of a
20 different document in step 234, then the process returns to step 226 to download initial portions of the different document.

Step 231 is implemented after step 227 if a particular embodiment of process 220 allows a user to
25 select a particular page of a new document to download and view, where no portion of that document has been previously downloaded. The header and range table of the document were downloaded in step 226. In step 231, the viewer connects to the host computer (if a new connection
30 is needed) and downloads the page offset table from the optimized document file. The page offset table is needed to determine the location of the requested page in the document file. In step 232, page information is determined for the document, similar to step 230
35 described above (and described with reference to Figure

- 54 -

11). The process then continues to step 236, as described below.

In step 236, the viewer requests a page offset from the finder so that the viewer can download the appropriate data. In step 238, the finder consults a page start offset table of the page information, which was created in step 230 or step 232 (as shown in Figure 11), to determine the first offset for the page requested by the user. Alternatively, the viewer can request a specific object rather than a page, and the finder can consult the table of first objects on a page (also created in step 230 or 232) to determine at which offset the object is located in the document file. The found first offset for the page is returned to the viewer in this step.

In step 240, the finder determines any additional ranges of data that are required to completely download and display the requested page, such as additional page contents and any shared objects for the page contents. The finder determines these additional ranges from the page information generated in step 230 or 232. If additional ranges of data are needed (as is typically the case), the page information indicates where the appropriate ranges are located in the optimized document file. In step 242, the viewer connects to the host computer (if a new connection is needed) to download the data of the requested page located at the first offset returned by the finder. In one embodiment, the viewer downloads a predetermined amount of data, such as one kilobyte (1K), at one connection.

In step 244, the finder requests any additional ranges of data for the requested page during the viewer connection. The finder may interleave particular shared objects in portions of the page content data in this request to provide a faster display of the page data for

- 55 -

the user. Step 244 is described in greater detail with respect to Figure 12. In addition, and optionally, the finder may delay requesting certain large objects such as fonts and images, as is described in greater detail with respect to Figures 13a, 13b, and 13c.

In step 246, the downloading of the additional ranges requested by the finder continues as a "background" process while the viewer processes the data already received. Once enough data is received by the viewer to display at least some page data, that data is displayed. The additional offset ranges may be stored in a local buffer before being requested by the viewer, such as in RAM and/or on disk. When the viewer requests data after the first block, it receives the data from the local cache rather than from the source file.

Other objects in the document file can be downloaded at later times when appropriate. For example, special objects such as page icons or bookmark objects can be downloaded if the user changes to an appropriate view in the viewer, scrolls to see more bookmark objects in a bookmark view, etc. Page tree objects can be downloaded if the viewer wishes to access such objects.

In step 247, the process places the requested page in a page cache implemented, for example, in local memory such as RAM 16 or on a hard disk or other storage device. This allows a previously-downloaded page to be quickly retrieved and displayed from the cache if the user should desire to view that page at a later time. Also in step 247, the process can place any appropriate shared objects that have been downloaded (along with their signatures, if any) into a "shared object cache" implemented, for example, in local memory such as RAM 16 or hard disk. These cached shared objects can be used when downloading and displaying page content of other documents as well as other pages of the current document that reference these

- 56 -

shared objects. The downloading of redundant copies of the cached shared objects from a document file thus can be avoided to speed up the downloading process. In one embodiment, the shared objects can be stored in the cache
5 over multiple downloads or even when computer system 10 is powered down (using battery backed RAM or other nonvolatile storage devices). The process is then complete at 248.

The downloading process of the present invention
10 allows a particular page desired by the user to be downloaded without downloading other pages in the document. This allows a speedier download since only the data for the page is downloaded, not the entire document file. In addition, new desired, randomly-accessible
15 pages of the document can be readily downloaded using the page offset table, giving the user the illusion that the entire document is easily available and accessible through simple commands such as "display next page" or "display page number X."

20 In prior art processes, when additional data for the page beyond the first 1K portion (or whatever sized portion is being used) is to be downloaded, an additional connection must be made for each additional 1K portion, which can cause delays in displaying the page data.
25 Connecting to a host computer over a network can be a lengthy process, since the host must be located over the network and appropriate handshaking signals must typically be sent and received, and the host computer/server may be busy. In the present invention,
30 however, the finder requests the additional ranges during the same connection that the viewer is requesting the first range, downloads the additional ranges as a background process, and stores the additional ranges in a local buffer. These additional ranges can be provided
35 directly from the buffer to the viewer when the viewer

- 57 -

requests additional 1K ranges. Thus, no additional connections to the host are required at a later time to download more data for the page. By avoiding multiple connections to the host in the present invention, the page data is downloaded and displayed much more quickly for the user. As an analogy, this one connection downloading process can be compared with going to a grocery store to buy grocery items, and buying all of the needed items on one trip rather than having to take multiple trips to buy items that were not bought on the first trip. The multiple trips are much more inefficient in the time involved than the single trip. In an alternative embodiment, described in reference to Figure 14, multiple ranges of data may be requested by the finder in one transaction for the page, avoiding the overhead of multiple transactions.

Figure 11 is a flow diagram illustrating step 230 of Figure 10, where the finder generates page information from the page offset table. The page information is to be used by the finder to request additional ranges of page data when the viewer is downloading page data. The process begins at 250, and, in step 252, a page start offset table is generated from data in the page offset table. The page start offset table includes the starting offset, e.g., in bytes, for the page contents of each page in the electronic document. More specifically, information in the page offset table such as the number of pages in the document, number of objects on a page, length of a page, etc., are used to create the page start offset table. Referring to Figure 11a, a page start offset table 264 is shown, where each entry 266 is the starting offset for each page's contents. The length of a page's contents (and shared objects) can be determined by taking the difference between the page's starting offset and the next page's starting offset, since all of

- 58 -

the contents of the page were arranged contiguously in the optimized file creation process of Figure 4.

In next step 254, a table that includes the object ID of the first object on each page is generated from the information in the page offset table. This table is generated from the number of objects on each page in the page offset table. This information can be decompressed from the page offset table using the size in bits to represent the number of objects on a page-- fourth entry in page offset table-- as is well known to those skilled in the art. Other information in the page offset table can be decompressed similarly. As shown in Figure 11a, the table 268 of the first object on each page includes an entry 270 for each page, where the object ID of the first object on the page is stored. The entries 270 of table 268 correspond to the order and number of starting offset entries 266 in table 264, so that a page can be indexed similarly in either table 264 or 268. Table 268 can be used to reference objects if the viewer asks for a page's starting offset based on an object ID rather than a page number. The page number for an object can be found in table 268, and the starting offset of the page then can be referenced in table 264.

In step 256, a shared object offset table is generated from the page offset table. The shared object offset table includes a starting offset, e.g., in bytes, for each shared object in the document. The shared object offsets may be stored in this table in the order of shared objects as stored in the document file. As shown in Figure 11a, shared object offset table 272 includes entries 274 for storing the shared object starting offsets in the document file. The shared object offset table can be generated from the information in the page offset table, such as the lengths of shared objects, number of shared objects, least size of a shared object,

- 59 -

and the location where the shared objects start (which can be calculated from the last page offset (from table 270) + last page length in the page offset table).

In next step 258, the sharing pages list 136 is generated from the page offset table information. This list 136 is stored in memory, such as RAM 16, of the client computer system 10, and includes substantially the same nodes as shown in Figure 6a. From list 136, the finder can determine which pages refer to which shared objects, the fraction of the page content in which the shared objects are referenced, and the sharing ID's of the shared objects so that the starting offsets for the shared objects can be referenced in shared object offset table 272. The list 136 can be generated from the shared object flag, fractions, and sharing ID's in the page offset table. The process is then complete as indicated at 260. In other embodiments, page information can be organized in different ways, or can be directly used from a page offset table.

Figure 12 is a flow diagram illustrating step 244 of Figure 10, in which the finder requests any additional ranges of page data for the requested page during the viewer's connection with the host computer. The process begins at 280. In step 282, the finder requests cross reference table data for the page contents object from the document file at the host computer for the requested page. The cross reference data allows the viewer to determine the offsets for objects other than the first object on the requested page. In next step 284, the process checks if interleaving should be performed on the downloaded page data. Interleaving, as described above, is the insertion of certain shared objects after portions (or "fractions") of page contents data that refers to those shared objects. This speeds up the display of portions of a page, since the shared objects required for

- 60 -

displaying a portion of page contents are downloaded right after that portion is downloaded. Alternatively, as described in greater detail with respect to Figures 13a, 13b, and 13c, the requesting of certain large
5 objects, such as embedded fonts and images, may be delayed until after requests for other, smaller page contents data have been made.

If no interleaving of page data is desired, the process continues to step 300, described below. If
10 interleaving of the page data is desired, the process continues to step 286, where the finder checks if there are more shared objects referenced by any remaining fractions of the page contents object in the document file by checking sharing pages list 136. If so, then
15 process continues to step 288, where the next shared object is examined from the appropriate object list 140 in the sharing pages list 136. If a cache is implemented, step 290 is performed, in which the process checks if the examined shared object is in the cache.
20 For example, if signatures are used, the finder can compare the signatures of objects in the cache with the signature of the examined shared object. If the examined shared object is already in the cache, then this object need not be downloaded during the current process. Thus,
25 the process returns to steps 286 and 288 to examine the next shared object in object list 140.

If the examined shared object is not in the cache, step 292 is performed, where the finder requests a range of page contents data from the start of the remaining
30 portion of page contents to the end of the fraction of page contents that references the examined shared object. The process determines that fraction by examining the fraction field 146 of the examined shared object in the object list 140 (which is the dividend of the fraction).
35 The divisor of the fraction is known from the page offset

- 61 -

table 68. An actual fraction size of page contents (in bytes) can be calculated by deriving an approximate page contents size from the page length and approximate page contents fraction (in the page offset table), and

5 dividing the approximate page contents size by the divisor. For example, if a fraction divisor of 8 is used, the next eighth of the page contents is requested by the finder in step 292 if the examined shared object is referenced by this eighth. If the examined shared

10 object is referenced by the sixth eighth of page contents, then all the page contents data (in the page contents object) up to and including the sixth eighth is requested by the finder. Thus, any fractions of page contents that do not reference a shared object in the

15 object list are combined until a shared object is found in the object list, at which time the request for the entire range of combined fractions is sent to the host computer.

The shared objects in an object list 140 are

20 provided (in the process of Figure 5) such that no shared object in the object list is duplicated due to being referenced multiple times on a single page. Thus, a shared object is only downloaded once for a page.

In an alternative embodiment, step 290 can be

25 omitted and only the next fraction of page contents can be requested in step 292 regardless of whether that next fraction references a shared object in the object list 140. If that next requested fraction does not include a reference to a "new" shared object (i.e., a shared object

30 that has not been previously downloaded), then step 294 (detailed below) can be skipped. In effect, this allows step 292 to be consecutively implemented multiple times until a new shared object is found and step 294 is implemented.

- 62 -

In next step 294, the finder requests ranges of the document file from the host computer for any non-cached shared objects referenced by the range of page contents data requested in step 292 (the finder also
5 requests appropriate cross reference table information for those shared objects). The finder requests the range of data for the shared object examined in step 290, and also requests ranges for any additional shared objects that are referenced by the requested page contents and
10 which are not in the shared object cache (the finder can compare objects in the cache for each such additional shared object). The finder can determine which shared objects were referenced in which fraction by examining the shared pages list 136, and the offsets for the shared
15 objects are known from the shared object offset table 272. For example, in the described embodiment, the finder can examine the fraction field 146 of nodes in object list 140 and request ranges for all shared objects having the same fraction number as the examined shared
20 object. The finder also requests any cross reference table information from the cross reference table 58 that references these shared objects. The process then returns to step 286.

It should be noted that, with shared image objects
25 placed at the end of each object list 140 in the sharing pages list 136 and forced to be in the last fraction of page contents, the image objects are forced to be downloaded after all page contents and other interleaved shared objects. Alternative embodiments for the display
30 of image objects are described in greater detail with respect to Figure 13c.

If no more shared objects are referenced by remaining fractions of the page contents object in step 286, then the process continues to step 296, where the
35 finder requests a range including all of the remaining

- 63 -

page contents data (if any unrequested page contents still remain in the document file). In next step 298, the finder then requests any non-contents shared objects for the page and the appropriate cross-reference table
5 information for those shared objects. As explained above, non-contents shared objects can include objects such as user-specified objects, resource objects which are not forced to be shared (such as procsets), etc., and which are shared. These non-contents shared objects can
10 be found in the object list 140 of the sharing pages list 136. The process is then complete as indicated at 302.

Under some conditions, interleaving may not be desirable. For example, when the page contents are very small, such as under 4 kilobytes, interleaving may
15 provide no significant display speed increase, as described above with reference to Figure 9. If no interleaving is desired, then the process continues from step 284 to step 300, wherein the finder requests the rest of the contiguous page contents data without
20 interleaving, and also requests the other objects stored after the page contents object that are needed for the display of the page contents object. The "other objects" include shared objects, any appropriate special objects, and any additional cross reference information needed for
25 the objects. For example, the shared objects can be requested in the order they are referenced on the page. In alternative embodiments, the shared objects can be requested first, followed by the page contents. The process is then complete at 302.

30 Turning to Figure 13a, a viewer program may optionally employ a method for progressive rendering of fonts to display text using a substitute font when the desired font (for example, an embedded font specified in an electronic document) is not immediately available.
35 The method for progressive rendering of fonts could be

- 64 -

invoked, for example, in connection with steps 242, 244, and 246, illustrated in Figure 10. Use of this method is advantageous, for example, in connection with embedded fonts in electronic documents downloaded over a network.

5 An embedded font need not be in the same file as the main part of the document, however; a document may refer to embedded fonts indirectly so that they can be obtained from another file or from a shared font resource or server. When a font reference is encountered, step 400,
10 the software must determine whether the desired font is already available to it, step 402. This would be the case, for example, if the font had already been stored in a cache (for example, in step 247, illustrated in Figure 10) or stored on a local data store, such as a hard disc
15 drive 28 (illustrated in Figure 1). If the font is available, step 404, naturally the software would use it.

On the other hand, step 406, if the desired font is not available, the software obtains a substitute font. A substitute font is appropriate to the extent it has and
20 metrics that are close to those of the desired font, so that the text drawn using the substitute font will appear in substantially the same place and form as text drawn using the desired font, which is done in steps 420 through 428. In one embodiment, a substitute font is
25 created based on a complete set of font metrics that are specific to the desired font. In PDF format documents, for example, a font descriptor object provides metrics including dimensional information such as ascent, descent, boundary block, height of capital letters,
30 italics angle, and width of vertical stems, and also including other descriptive information such as whether this is a serif or sans serif font, whether it is a fixed width font, whether it is an all caps font, and whether it is mixed height all caps font. A second PDF object,
35 the width array object, completes the description of font

- 65 -

metrics: it provides width information for each letter in the font. Using such metrics contained in the document, or metrics obtained from a database or a data file on computer system 10 or on a server on a network, or, in
5 the absence of better information, default metrics, an appropriate substitute font may be created by using a font manager software module, such as the Font Chameleon product available from Ares Software Corporation of Foster City, California, or the Infinifont product
10 available from ElseWare Corporation of Seattle, Washington. The Acrobat reader product uses the Adobe Multiple Master font technology to create substitute fonts. Alternatively, the available font metrics, if any, for the desired font may be used to adopt a font
15 from among the font resources available on, or readily available to, computer system 10. When a substitute font has been obtained, whether by creation or adoption, it is used, step 408, to render the corresponding text; and the software continues processing the document, step 410.
20 For situations in which the desired font has characters that are not generally available in fonts -- such as a ligature fi character, for example -- or the desired font is a set of graphics -- such as chess pieces or icons, for example -- the substitute font may have
25 place-holding symbols, or even consist entirely of just one place-holding symbol (such as an empty box), that is, or are, displayed in accordance with the available font metrics.

When the desired font becomes available, step 420,
30 because it has been downloaded as embedded font in the electronic document, for example, the affected portions of the display are redrawn, steps 422 through 428. (If more than one desired font had been replaced by a substitute font, all of the desired fonts could be
35 processed together in these steps.) First, the bounding

- 66 -

rectangle or rectangles of any text rendered using the substitute font are determined, step 422. These rectangles surround the areas used drawing with the substitute font and the areas to be used drawing with the desired font. Next, step 424, an off-screen bitmap buffer is created with the boundaries corresponding to the intersection of the bounding rectangle or rectangles, limited to the visible portion of the page. The off-screen buffer is created to have the same pixel depth and color characteristics as the display. Then, in step 426, the available parts of the page that intersect with the off-screen bitmap buffer are redrawn into that buffer using the desired font or fonts, which are now available. (Some parts of the page, such as a large image, for example, may not be available, and the process should proceed without them, as will be described in reference to Figure 13c.) When the off-screen bitmap buffer is complete, it is drawn onto the display, step 428. If the display is generated, as is conventional, from a display buffer, this is done by copying the off-screen bitmap buffer into the corresponding area of the display buffer. In an alternative embodiment, the unit of display is some multiple (or fraction) of a page, and the processing of embedded fonts is delayed until the entire unit of display has been drawn, allowing the user (reader) of an electronic document to see text and graphics on all of the visible region as soon as possible. In a further alternative embodiment, drawing into an off-screen bitmap buffer is delayed until all desired fonts are available, which can reduce the time spent rendering off screen. In a further alternative embodiment, multiple small rectangles are joined to form one or more larger rectangles when creating off-screen bitmap buffers, which also can simplify the processing of the off-screen bitmap buffer. In a further alternative embodiment, the method

- 67 -

is applied only to some of the universe of possible desired fonts, so that foreign language fonts, for example, are excluded. The Acrobat reader product, for example, does not apply progressive rendering if the
5 desired font has characters not found in the Adobe standard character set; nor does it apply progressive rendering in the absence of font-specific font metrics.

It should be noted that to take maximum advantage of the methods for progressive rendering of fonts, one
10 should control the location of embedded fonts in the process of reading or downloading a document, so that embedded fonts that are subject to substitution are read or downloaded after the main text and graphics. However, if extraction of embedded fonts is time consuming, it may
15 be advantageous to use the method even if the embedded fonts are located before, or are intermingled with, the text and the graphics. It is also advantageous, in connection with this method, to be able to download embedded fonts with a background process while
20 maintaining a user interface during the initial display of the document. Finally, it will be understood that if detailed and complete font metrics are available, the desired font can be emulated very closely by the substitute font and the layout of text on the display can
25 be preserved and user distraction can be minimized when the display is updated. User distraction is minimized because the update of the display will be a ripple of subtle character shape changes rather than more significant changes to the flow of text in the document.
30 If the descriptive information is less complete or specific, or if default font metrics have to be used, more significant reformatting of the document will occur, including, possibly, significant reposition of the text in the final display. Further information on techniques
35 of font substitution may be found in commonly-assigned

- 68 -

United States Patent No. 5,185,818 to Warnock for Method of Substituting Fonts and Displaying Characters, the disclosure of which is incorporated herein by this reference.

5 Turning to Figure 13b, another aspect of optimizing user interaction with electronic documents in on-line environments is to ensure that interactive document elements, such as hypertext links, are active in the user interface as soon as possible. The term
10 "annotation" will be used herein in the sense it has in the context of PDF documents: the term "annotation" includes interactive elements that are associated with a page but not properly part of the printed page itself, such as hypertext links, icons selectable to invoke a
15 multimedia presentation, and so on, which respond to user input from devices such as a keyboard and mouse. Making annotations responsive to user input before they are rendered is useful to users that frequently visit the same document on-line. Such a user may know that a
20 particular region of the display has an annotation and choose to click there before any graphics for the annotation are drawn. Thus, in an optional method for early activation of active elements, when an active element to be displayed is encountered while processing
25 the document for display, step 440, the software first (or in the foreground) identifies the selectable boundary of the active element, step 442; instructs the cursor display process to change the appearance of the cursor appropriately when the cursor is displayed within the
30 selection boundary, step 444; and enables response to selection activity by the user, such as clicking or entry of keyboard commands, step 446. Second (or in a background process or otherwise in the normal course of displaying the document), the display associated with the
35 active element is rendered, step 448.

- 69 -

Turning to Figure 13c, a viewer program may optionally employ a method for deferring (delaying the downloading and/or display of) large objects. This method provides useful optimization of the display of electronic documents, particularly in on-line environments, by displaying information and text, graphics, annotations, and other quickly-displayed elements before large and/or complex objects, such as bitmap images, are downloaded and rendered. Electronic documents in formats such as PDF are capable of representing complex relationships between text, graphics, images, and annotations. Page contents can be overlapping and can be specified in the electronic document to have a specific display order. In the method for deferring large objects, if the electronic document specifies that a segment of text, for example, should appear on top of an image, for example, the text will be drawn first, allowing the using to be interacting with the portions of the electronic document as soon as they become available. Thus, when a large object is encountered, step 460, the rendering of the object is deferred, step 462, and the quickly-displayed elements, such as text and graphics, are drawn as soon as they become available, step 464. In electronic document formats, such as PDF, that precisely describe the exact location of every object on a page, the delayed rendering of some objects does not cause other objects to shift their locations, minimizing user disturbance as page elements are progressively displayed.

Where text that has already been drawn is specified as appearing on top of an image that has not yet been drawn, to continue the example, an off-screen buffer is created, step 466, into which are rendered the deferred object and all other drawing elements in their proper order, to the extent they appear within the

- 70 -

boundaries of the off-screen buffer, step 468, but without waiting for any objects (such as images yet to be downloaded) that are not then available. The boundary rectangle of the off-screen bitmap buffer corresponds to the boundary of the large object, limited to the visible portion of the page. The off-screen bitmap buffer is created to have the same pixel depth and color characteristics as the display. When the off-screen buffer is completed, it is drawn onto the display, step 470. In an alternative embodiment, for a page oriented viewer, if more than one page is visible at once, the rendering of large objects is deferred until all visible pages have been drawn, to allow the reader of the document to see the quickly-displayed content on all of the visible pages as soon as possible. In a further alternative embodiment, parts of an image are rendered into the image's off-screen buffer as soon as the parts become available, and then displayed, rather than waiting for the entire image to be downloaded. This may readily be done, for example, in image object formats that organize image information in bands. In a further alternative embodiment, the processing of several large objects is combined in an off-screen buffer defined by the union of their boundaries, which can reduce the time spent in rendering off screen.

LINEARIZED OPTIMIZED DOCUMENT LAYOUT

Turning to Figure 14, a linearized document layout is a refinement of the optimal format that has already been described. The linearized layout (or format) allows a view to achieve efficient incremental access in an environment having the characteristics of a low speed connection to a World Wide Web site. Among the advantages of the linearized format are that, when an electronic document is opened, display of a first page

- 71 -

occurs quickly. The first page to be viewed can be an arbitrary page of the document; it is not necessarily page zero, although would often be the most common choice. Other advantages are shared with other

5 alternative embodiments of the optimized electronic document format, which have already been described, including the quick display of pages after they are requested, the incremental display of page contents (particularly for pages delivered over a slow channel)

10 with the most useful data being displayed first, and the enablement of user interaction, such as following a hypertext link, before an entire page has been received and displayed. The optimized formats, including the linear format, achieve these advantages for electronic

15 documents of arbitrary size, with the total number of pages in the document having little or no effect on user-perceived performance in viewing any particular page.

The World Wide Web environment has characteristics that affect the performance of clients reading, and

20 servers providing, electronic documents. The access protocol (HTTP) is a transaction consisting of a request and a response. After a transaction has completed, obtaining more data requires a new request-response transaction, and the connection between client (viewer)

25 and server does not ordinarily persist beyond the end of a transaction, although some implementations may attempt to cache an open connection in order to expedite subsequent transactions with the same server. Round-trip delay can be significant: a transaction can take up to

30 several seconds, independent of the amount of data requested. And finally, the data rate may be limited: a typical bottleneck is a 14.4 kilobaud or 28.8 kilobaud modem link between the client and an Internet service provider. Other wide-area network architectures

35 generally share these characteristics. Even CD-ROMs

- 72 -

share some of these characteristics, since they have relatively slow seek times and limited data rates compared with magnetic media.

In a soon-to-be-supported extension to the HTTP
5 protocol, a client can request retrieval of portions of a document by specifying one or more offset-length byte ranges as part of the transaction request (that is, as part of the URL). Each range can be relative to either the beginning or the end of the file. The client can
10 specify any number of ranges in the request, and the response will consist of multiple blocks, each properly tagged. In some environments, including some World Wide Web environments, the client can initiate multiple concurrent transactions in an attempt to obtain multiple
15 responses in parallel. This is commonly done, for instance, to retrieve in-line images referenced from a HTML document. However, because multiple concurrent transactions appear to be less than optimal for PDF format documents in some important environments, the
20 linearized format is designed so that good performance is achieved under the constraint that only one transaction is active at a time. For that reason, the linearized format provides the client sufficient information to determine the byte range for every object required to
25 display a given page so that the client can specify the appropriate byte ranges in a single request.

Turning now to its implementation in a PDF format document, the linearized layout begins with a conventional PDF header 480 and trailer 504. The header
30 is followed by an object 482 containing linearization parameters, which include the format version identifier; the length of the entire file; an offset-length pointer to the primary hint table stream object 486 (described below); an optional offset-length pointer to an optional
35 overflow hint table stream object 500, if any; an object

- 73 -

number of the first page's page object; and, optionally, the page number of the first page to be displayed, which by default is page zero.

The first page cross-reference table 484 is a cross-reference table (such as was described in reference to tables 58 in Figures 3a and 3b) for all the first page's objects and all document-level objects appearing before the first page. This is a cross-reference table for all of the objects needed to display the first page of the document, including the parent page object for the first page, which has no attributes inherited from any other page's page object. The objects tabulated in table 484 include all objects that the parent page object refers to, to any arbitrary depth, if they affect the display of the page. It also contains entries for the dictionary of hint tables at the beginning of the primary hint table stream object 486. The first page cross-reference table 484 is a valid cross-reference table according to the PDF specification, although its position in the document file is unconventional. Its trailer portion gives the offset to the main cross-reference table 502 near the end of the file, as well as any cross-reference table attributes required to display the document.

The primary hint table stream object 486 may either precede or follow the first page objects 492. In an alternative embodiment, the data contents of the primary hint table stream object 486 are stored in one or multiple secondary files, residing optionally on separate host computers. A stream object is a type of object defined by the PDF format: it is a sequence of binary bits that may have compression associated with it. Each hint table consists of a portion of the stream object 486, beginning at the position in the stream indicated by the dictionary. It is expected that each table will

- 74 -

encode the required information as compactly as possible, because tables in the document need not be designed for random access, and that the client will read and decode the tables once and retain the information as long as the document remains open. The hint table stream object 486 includes a dictionary showing the positions of the hint tables in the object. These positions are relative to the beginning of the stream data.

A hint table encodes the position of a group or class of objects in the document file. The representation is either explicit, an offset from the beginning of the file, or implicit, from the cumulative lengths of preceding objects. In either case, the resulting positions are interpreted as if the hint table stream object itself were not present. This is so because the length of a hint table stream object is in general not known until after it has been generated and information in the hint table should not depend on knowing that length in advance. If an overflow hint table stream object 500 exists, obtaining it requires issuing an additional transaction. However, providing for an overflow object allows a linearizer program to write a linearized file with space reserved for the primary hint table stream object 486 of an estimated size, and then to go back to fill in the hint tables. If the estimate is too small, the linearizer program can append an overflow object containing the remaining hint table data, which allows writing the file in one pass, which may be advantageous in some circumstances. If there is an optional overflow hint table stream object, the contents of the two stream objects are to be concatenated and treated as if they were a single unbroken stream object.

Hint tables are not used in processing the objects of the first page, so their position relative to the

- 75 -

first page is a matter of choice. If the hint tables come before the first page objects, the client can abort the initial transaction quickly in the event that the document's selection of the first page to be displayed is not the same as the client's selection. On the other hand, if the hint tables follow the first page objects, displaying the document's selection of the first page is accomplished more quickly, while opening an arbitrary page is delayed by the time required to receive the document's first page. When an electronic document file is linearized, the linearization process may accept, as a user option, a decision whether to favor opening at a first page or opening at an arbitrary page.

Turning to Figures 15a and 15b, the hint tables of object 486 include a page offset hint table 488 and a shared object table 490, which two tables perform the same function as range table 64, shown in, and described in the context of, Figure 3b. The page offset table 488 gives, for each page, the information required to locate that page. Additionally, for each page except the first, it enumerates all shared objects that the page references, directly or indirectly. The shared object table 490 gives the information required to locate shared objects. In the linearized format, shared objects can be physically located in either of two places. Objects that are referenced from the first page are co-located with the first page objects. All other shared objects are located in the shared objects section 498. A single entry in the shared object table 490 can describe the group of adjacent objects, if only the first object in the group is referenced from outside the group. The page offset table 488 refers to an entry in the shared object table 490 by a simple index that is its sequential position in the table.

- 76 -

Other hint tables provide information for document elements that relate to the document as a whole. For example, the bookmark hint table allows the client to find bookmarks; and the thread hint table allows the client to find all the beads in thread of beads that defines an article. As in a newspaper, where an article may extend across several pages, and "article" here is a list of "beads", where each bead is a rectangle on a particular page, in which rectangle a portion of the article's text or illustrations may be found. With the thread hint table, the client (viewer) may request all of the objects required to display the entire article in one transaction.

Thus, hint tables provide indexing information that enables the client to construct a single request for all the objects required to display any page of the document or to retrieve certain other information efficiently. Hint tables may also contain information to optimize access to application-specific information by plug-ins.

Hint tables are not logically part of the information content of the document; they can be derived from the document. When the document is regenerated, the hint table stream objects 486 and 500 would not be part of the document unless they were specially generated. Any action that changes the document -- for instance appending an incremental update -- may invalidate some or all of the hint tables. The resulting document file would still be valid as a PDF file, but not necessarily as a linearized file.

Turning to Figure 15a, a few of the elements of page offset hint table 488 still need to be described. The "number of objects in page" is a value, which when added to the "least number of objects in a page", given in the header, give the actual number of objects in the

- 77 -

page. The first object of the first page has an object number given in the linearization parameters object 482. The first object of the second page has an object number of 1. Object numbers for subsequent pages can be
5 determined by accumulated the number of objects in all previous pages.

The "page length in bytes" is a value which, when added to the "least length of page in bytes", given in the header, gives the total length of the page in bytes.
10 The location of the first object of the first page can be determined from the cross-reference table entry for that object. The location of subsequent pages can be determined by accumulating the length of all previous pages. The "number of bytes from start of page to start
15 of contents stream" is a value which, when added to the "least start of contents offset", given in the header, gives the offset and bytes of the content stream object relative to the beginning of the page. The "length of contents in bytes" is a value which, when added to the
20 "least contents length", given in the header, is the length of the contents stream object in bytes, including object overhead.

The page offset hint table 488 includes, for each shared object referenced from each page, a "shared object
25 identifier" and "fraction giving position in contents of first reference". The former is an index into shared object hint table 490. The latter indicates where in the page's contents data the shared object is first referenced. As has been described, this is interpreted
30 as the numerator of a fraction, whose denominator is specified in the page offset hint table header. The numerator can take on additional values, nominally indicating fractions past the end of the contents data, to indicate that the shared object is not referenced from
35 the contents, but is needed by annotations or other

- 78 -

objects that are drawn after the contents. Different values may be used to designate when the shared object is needed with respect to the types of non-shared objects that are at the end of the page.

- 5 Shared object hint table 490 is illustrated in Figure 15b. The page offset hint table 488 refers to an entry in the shared object hint table 490 by a simple index that is the entry's sequence in the table. Following the header information, which is self-
- 10 explanatory, there are two sequences of shared object group entries: the ones for objects 492 located in the first page are followed by the ones for objects located in the shared objects section 498. The entries have the same format in both cases. For convenience of
- 15 representation, the first page (492 objects) may be treated as if it consisted entirely of shared objects. That is, the first entry of the shared objects table refers to the beginning of the first page and has an object count and length that span all of the initial non-
- 20 shared objects. The next entry refers to a group of shared objects in the first page objects 492. Subsequent entries span additional groups of either shared or non-shared objects consecutively, until all shared objects in the first page have been enumerated.
- 25 In the information that appears for each shared object group, the "number of objects in group" is a value that is one less than the actual number of objects in the group. The object number first object of the first page is given in the linearization parameters object 482 at
- 30 the beginning of the document. Object numbers for subsequent entries can be determined by accumulating the number of objects in all previous entries, until all shared objects in the first page have been enumerated. Following that, the first object in the shared objects
- 35 section 498 has a number that can be obtained from the

- 79 -

shared object table header, "object number of first
objects, in shared objects section". The "total length
of objects in bytes" is a value which, when added to the
"least length of a shared object group in bytes", given
5 in the header, gives the total length of the object group
in bytes. The location of the first object of the first
page is given in the page offset table header
information. The locations of subsequent object groups
can be determined by accumulating the lengths of all
10 previous object groups until all shared objects in the
first page have been enumerated. Following that, the
location of the first object in the shared objects
section 498 can be obtained from the header. The
"signature present flag" indicates the presence or
15 absence of a signature. The "optional signature" when
present is, in one embodiment, a 16-byte MD5 hash
intended to identify uniquely the resource that the group
of objects represents. This enables the client to
substitute a locally cached copy of the resource instead
20 of reading it from the document.

Turning to Figure 15c, each entry in the thumbnail
hints table 506 describes the thumbnail for a single
page. The pages are considered in page number order,
starting at page zero even if page zero is not the first
25 page to be displayed. Thumbnails may exist for some but
not all pages. The header for this table is self-
explanatory. In the entries, the "count of preceding
pages lacking thumbnails" indicates how many pages
without thumbnails lie between the previous entry's page
30 in this one. If all pages have thumbnails, the value of
this field is always zero and the value of "bits needed
to represent count of thumbnail-less pages" in the header
can be zero. The "length of thumbnail object in bytes"
is a value which, when added to the "least length of
35 thumbnail object in bytes", given in the header, gives a

- 80 -

total length of the thumbnail object. The "for each shared object: shared object identifier" is an index into the shared object hint table 490.

Turning to Figure 15d, a format for generic hint tables 508 is illustrated. Certain categories of objects are associated with the document as a whole rather than with individual pages. It is sometimes useful to provide hints for accessing such objects efficiently. For each category of hints that is supported with a hint table, there is an entry in the hint table stream object 486 giving the starting position of the corresponding hint table within the stream. The illustrated format of generic hint table 508 provides a generic representation for such hints. This representation is useful for standard categories of objects, such as outlines, threads, and named destinations. It may also be useful for application-specific objects accessed by plug-ins. The generic hint table 508 describes one or more groups of objects that are located together in the document. For each group, the hints contain sufficient information to enable the client (document reader process) to construct a request for all objects in the group, including any shared objects that the objects in the group may reference. When there is single group, the hint table refers to all the objects in the category, which are to be accessed at the same time. When there are multiple groups, each group is identified by a simple index that is its sequential position in the table, counting from zero. What the object groupings signify depends on the object category.

The header information for a generic hint table 508 is self-explanatory. In entries for each object group, the "total length of objects in bytes" is a value which, when added to the "least length of an object group in bytes", given in the header, gives the total length of

- 81 -

the object group. The "for each shared object referenced from group: shared object identifier" is an index into the shared object hint table 490. The number of bits needed to represent this identifier is given in the header of the page offset hint table 488.

In the first page object section 492 are all the objects, including shared objects, required to display the first page of the document. The following ordering of objects is useful for providing early user interaction an incremental display of first page data as it arrives. First, annotation objects to a depth sufficient to allow the annotations to be activated; information required to draw the annotations can be deferred, since annotations are always drawn on top of (hence after) other contents.

After first page object section 492 are sections containing, in sequence, the non-shared objects 494 for the next page through the non-shared objects 496 for the last page. For each page other than the first page to be displayed, the objects required to display the page are grouped together, except for resources and other objects that are shared with other pages. Shared objects are located in the shared objects section 498. The starting file offset and length of any page can be determined from the hint tables. The order of objects in non-first pages, like the order in the first page, should facilitate early user interaction and incremental display of page data as it arrives. Generally, in the linearized format, there will be little benefit from interleaving contents with resources because most resources other than images -- in fonts in particular -- are shared among all of the pages and therefore reside in the shared objects section 498. Image objects are usually not shared, but they should appear after other page objects because the rendering of images can be deferred, as has been described in reference to Figure 13c.

- 82 -

The shared objects section 498 contains objects that are referenced from more than one page and that are not referenced (directly or indirectly) from the first page. The shared object table 490 contains an index of 5 these objects. The order of the objects in the shared object section 498 is essentially arbitrary. However, it is desirable that where a resource consists of a multiple-level structure, all components of the structure are grouped together. Thus, if only the top-level object 10 is referenced from outside the group, the group can be described by a single entry in the shared object table 490, minimizing the size of that table.

The main cross-reference table 502 is the cross-reference table for all objects in the file other than 15 those listed in the first page cross-reference table 484.

An electronic document may also have other objects (not shown) that are part of the document but are not required for displaying pages. Such objects should be divided into functional categories, and objects within 20 each category grouped together, so that each may have a corresponding hint table to provide the information required for efficient access by a client. The linearized format allows for additional hint tables for application-specific data accessed by plug-ins. Such 25 additional hint tables may have a generic format or the format of the hint table can be private to the application.

From the foregoing, it will be understood that the linearized format allows for the efficient retrieval and 30 display of electronic documents. Thus, when a document is initially accessed, the client can issue a request to retrieve the entire file starting at the beginning. The data for the first page to be display will thus quickly appear and be available. Like the objects for the first 35 page, the primary hint table stream object will also be

- 83 -

retrieved and part of the initial sequential read of the file. It is expected that the client will interpret and retain all of the information in hint tables. When the hint tables and first page objects have been read, the client may decide to continue reading the remainder of the document sequentially, or the client may decide to abort the initial transaction and access subsequent pages using separate transactions requesting byte ranges. As soon as the hint tables have been received, the client has sufficient information to request retrieval of any page of the document, given its page number.

The ordering of objects and pages and the use of hint tables in the linearized format allows for progressive update of the display and early opportunities for user interaction when data arrives slowly. To this end, the client (viewer) should recognize whether objects referenced on a page have arrived and, where possible, adapt the order in which it acts on objects to the object's nature and availability. One such order of action is the following, the elements of which have been described in reference to Figures 13a, 13b, and 13c: first, activate annotations without drawing them; then draw contents but defer unavailable images and use substitute fonts for unavailable fonts; then draw annotations; then draw images together with anything that overlaps them; and then re-draw text using desired fonts, together with anything that overlaps the text. As has been described, the late drawing of images and the re-drawing of text may be done using an off-screen buffer; however, these may also be drawn directly into the buffer from which the display is generated.

EXAMPLE OF DOWNLOADING A REQUESTED PAGE

In an example of the processes of Figures 10 and 12, the first page of an electronic document is

- 84 -

downloaded from an optimized document file of the present invention and displayed in steps 224, 226, 228, and 230 of Figure 10. The user then requests to display page 9 in the viewer at step 234 of Figure 10. In this example, 5 the data needed to display page 9 occurs in the byte ranges of 25000-29000, 112000-113000, and 200000-202000 in the optimized document file, where the first range is the page contents data, and the second two ranges are shared objects referred to by the page contents. The 10 finder finds the cross reference table and offset 25000 for page 9 in step 238 and provides the offset to the viewer. The finder determines the additional ranges at step 240. At step 242, the viewer connects to the source computer and downloads a predetermined amount of data, 15 such as 1K. The finder, meanwhile, has determined that the page contents actually continues to byte 29000. If no interleaving is desired, the finder additionally requests byte range 26000-29000 to be downloaded in step 300 of Figure 12. Also, the finder has determined by 20 examining sharing pages list 136 that page 9 has two shared objects with sharing ID's of 1 and 2. The finder thus consults the shared object offset table 272 and requests the corresponding shared object byte ranges 112000-113000 and 200000-202000 to be downloaded at the 25 same connection in step 300 of Figure 12. Alternatively, as described in reference to the linearized optimized format illustrated in Figure 14, all of the byte ranges may be requested in one transaction, so that steps 238, 240, 242, 244, and 246 of Figure 10 operate as one step. 30 If interleaving is implemented, then in step 288 of Figure 12, the finder examines the next shared object in the object list, which has a sharing ID of 1 and a fraction number of 1. In step 292, the finder requests page contents from the beginning of the page to first 35 fraction. Thus, assuming the fraction is 0/8, each

- 85 -

fraction is $4000 / 8 = 500$ bytes. Since the first fraction references the examined shared object, the finder requests a range of 25000-25500. In step 294, the finder then requests ranges for any shared objects in the first fraction; there is only one in the object list. The finder thus requests the range of 112000-113000 for this shared object as determined from the shared object offset table 272. Since the 5th eighth of the page contents includes a reference to the only other shared object on the page (sharing ID of 2), the process returns to step 286 after step 294. In the next iteration, the finder examines the next shared object in the object list, which has a sharing ID of 2 and is in fraction number 5. The next four eighths of page contents are thus retrieved in step 292, up to and including the 5th eighth that includes the examined shared object (or, alternatively, step 292 can be implemented multiple (5) times by requesting one fraction each time through the loop, and skipping step 294 until the shared object reference is found). Thus, a range of page contents from bytes 25501-27500 of the document file is requested (four fractions = 2000 bytes). Then, in step 294, the shared object referenced by the 5th fraction at range 200000-202000 (and its cross reference information) is requested. Since no further shared objects are referenced in the page contents, the finder requests the remaining portion of page contents in step 296, which has a byte range of 27501-29000 (and its cross reference information). All these byte requests are performed during a single connection to the host computer. Alternatively, as described in reference to the linearized optimized format illustrated in Figure 14, all of the byte ranges may be requested in one transaction.

Thus, the process of the present invention allows data for an entire requested page to be read from a

- 86 -

source computer with only one connection to (or, alternatively, one transaction with) the source, after the first page or the initial set of tables has been downloaded. The page data can thus be downloaded and
5 displayed more quickly, with less waiting time for the user. In the interleaving process, shared objects can be downloaded right after portions of page content that reference those shared objects. This allows portions of a page to be displayed to the user immediately, without
10 having to wait for the shared objects to be downloaded.

While this invention has been described in terms of several particular embodiments, it is contemplated that alterations, modifications and permutations thereof will become apparent to those skilled in the art upon a
15 reading of the specification and study of the drawings. For example, the present invention is described as being used for portable electronic documents, such as PDF documents. However, other files or collections of data which, for example, include disjointed objects/data,
20 and/or shared objects/data are well suited to be optimized and downloaded by the present invention. In addition, the shared object interleaving of the described invention is not necessary to provide an optimized file for faster downloading. Also, many of the steps or
25 processes described in the described embodiments are specific to a described embodiment, and can be changed or omitted in other embodiments. For example, the use of the finder and viewer in the downloading process of Figure 10 can be changed to suit a particular embodiment.
30 Furthermore, certain terminology has been used for the purposes of descriptive clarity, and not to limit the present invention. For example, it is not intended that the term "table" be read narrowly to include only data structures having a conventional tabular structure;
35 rather, the term should encompass all forms of data

- 87 -

structure or structures that carry the required information. It is therefore intended that the following appended claims include all such alterations, modifications and permutations as fall within the true
5 spirit and scope of the present invention.

What is claimed is:

- 88 -

1. A method for reading a specific page of an electronic document, the method comprising:
 - reading page offset information early during reading of the document; and
 - 5 using the page offset information to locate the contents of the specific page, whereby the specific page is read without the necessity of reading other pages in the document.
2. The method of claim 1 where
 - 10 the reading of the document comprises downloading the electronic document across a computer network; and
 - the page offset information comprises a page offset hint table
- 15 3. The method of claim 1 further comprising:
 - reading a shared object hint table before reading a second page of the document.
4. The method of claim 1 further comprising:
 - 20 reading a bookmark hint table before reading of a second page of the document.
5. The method of claim 1 further comprising:
 - reading an article thread hint table before
 - reading of a second page of the document; and
 - reading a thumbnail hint table before the reading
 - 25 of a second page of the document.
6. The method of claim 1 wherein each category of objects associated with the document as a whole has a corresponding hint table.
7. The method of claim 2 wherein the page offset
- 30 information is read before the downloading of more than one page of the document has been completed.
8. The method of claim 2 further comprising a step of displaying the specific page requested by the user on an output display device.

- 89 -

9. The method of claim 8 wherein a portion of the specific page is downloaded by a viewer and all remaining portions of the specific page are determined by a finder process using the page offset table and downloaded during
5 one connection with the host computer.

10. The method of claim 2 wherein the specific page includes page contents and shared objects, where the shared objects are downloaded interleaved between portions of the page contents.

10 11. A method for providing an optimized document file, the optimized document file including a plurality of pages, the method comprising the steps of:

providing document information in the optimized document file, said document information
15 including page content information that describes individual pages of the optimized document file, wherein said page content information for said individual pages is provided contiguously in the optimized
20 document file; and

providing a page offset table in the optimized document file, the page offset table
including page offset information to be used
to locate the document information for
25 individual pages of the document file.

12. The method of claim 11 further comprising the step of providing first page offset information for a first page of the optimized document file, the first page offset information describing the locations of all
30 portions of the first page in the optimized document file, the first page offset information being provided separate from said page offset information.

13. The method of claim 12 further comprising providing special objects not needed for the display of the first
35 page of the document after the shared objects in the

- 90 -

optimized document file.

14. Apparatus comprising a computer-readable storage medium tangibly embodying program instructions comprising instructions forming a display process to:

- 5 display on the display screen a page-based document stored on a host computer, connect with the host to download page offset information located at a predetermined location in the page-based document;
- 10 download a specific page of the page-based document requested by a user without the necessity of downloading other pages in the document; and
- display the downloaded page on the display screen.

15 15. The apparatus of claim 14 wherein the storage medium further comprises instructions to:

- download a first portion of page content on a page of the page-based document, the portion of page content including a reference to a
- 20 shared object;
- download the shared object referenced by the first portion of the page; and
- download a second portion of page content of the page of the page-based document.

25 16. The apparatus of claim 15 wherein the shared object is downloaded from the page-based document.

17. The apparatus of claim 15 wherein the instructions further comprise instructions to:

- derive the locations of the first portion of page
- 30 content, the second portion of page content, and the shared object in the page-based document utilizing a page offset table downloaded from the page-based document.

18. The apparatus of claim 15 wherein the instructions

35 further comprise instructions to:

- 91 -

display the first portion and the second portion of page content on an output display device, wherein use of the shared object is needed to display the first portion of page content.

5 19. The apparatus of claim 15 wherein the page content includes text to be displayed, and the shared object is a font object needed to display the text.

20. The apparatus of claim 14 wherein the computer-readable storage medium is a program instruction
10 store, the apparatus further comprising:

a digital processor coupled to the program instruction store;

a display screen coupled to the digital processor; and

15 instructions stored in the program store forming a finder process, comprising instructions to: use the page offset information to provide a location of the specific page in the document to the display process so that
20 the display process can download the specific page.

21. The apparatus of claim 20 wherein when the display process is downloading a portion of a specific page during a connection with the host, the finder process
25 additionally requests additional portions of the specific page, and the additional portions of the page are downloaded to the display process to be displayed.

22. The apparatus of claim 20 wherein the page offset information is read before the downloading of more than
30 one page of the document has been completed.

23. The apparatus of claim 20 wherein the finder can derive the beginning offsets of each page of the document from the page offset information.

24. The apparatus of claim 21 wherein the display
35 process is coupled to a viewer process in the program

- 92 -

memory comprising instructions to select and display pages of the page-based document on the display screen.

25. A method for displaying on a display device of a computer an electronic document having an object and text
5 specified as having to be displayed as if drawn on top of the object, the method comprising:

deferring displaying the object in favor of
displaying the text;

displaying the text;

10 then displaying the portion of the object that is specified to appear as if drawn underneath the text; and

then displaying the text again, whereby the text is displayed as if drawn on top of the
15 object.

26. The method of claim 25, further comprising:

drawing the portion of the object that is specified to appear as if drawn underneath the text (the underneath portion) into a
20 buffer the contents of which are not displayed (the off-screen buffer);

drawing the text into the off-screen buffer after the underneath portion has been drawn into the off-screen buffer; and then

25 displaying the contents of the off-screen buffer, whereby the text is displayed as if drawn on top of the object.

27. The method of claim 26 wherein

the object is a large object or a bitmap image;

30 the text includes all of the text specified in the document to appear displayed on top of the large object; and

the underneath portion that is drawn into the off-screen buffer is substantially the entire
35 object.

- 93 -

28. The method of claim 26 further comprising:
downloading the document to the computer from an
other computer across a network; and
performing the step of displaying the text without
5 waiting for that part of the document
containing the object to be downloaded.
29. The method of claim 26 further comprising:
displaying a cursor that moves with a mouse or
pointing device;
10 changing the appearance of the cursor to indicate
when it is located where an annotation in a
PDF format electronic document will be
displayed; and
making the annotation in a PDF format electronic
15 document responsive to input from the user
without waiting for the annotation to be
displayed.
30. A method for displaying on a display device coupled
to a computer an electronic document having text
20 specified as having to be drawn using a font that is not
on the computer (the desired font), the method
comprising:
initially displaying the text using a substitute
font different from the desired font;
25 obtaining the desired font on the computer; and
then
redisplaying using the desired font that area in
which the substitute font was used in the
step of initially displaying the text.
- 30 31. The method of claim 30, wherein the document or the
computer has font description metrics for the desired
font, the method further comprising:
using the font description metrics for the desired
font to create the substitute font.

- 94 -

32. The method of claim 30, further comprising:
adopting a font from a font resource on a local
data store as the substitute font.

1/24





Figure 1

2/24



93

How to use this online guide

21.42

46

42

42

45

Figure 2a

3/24

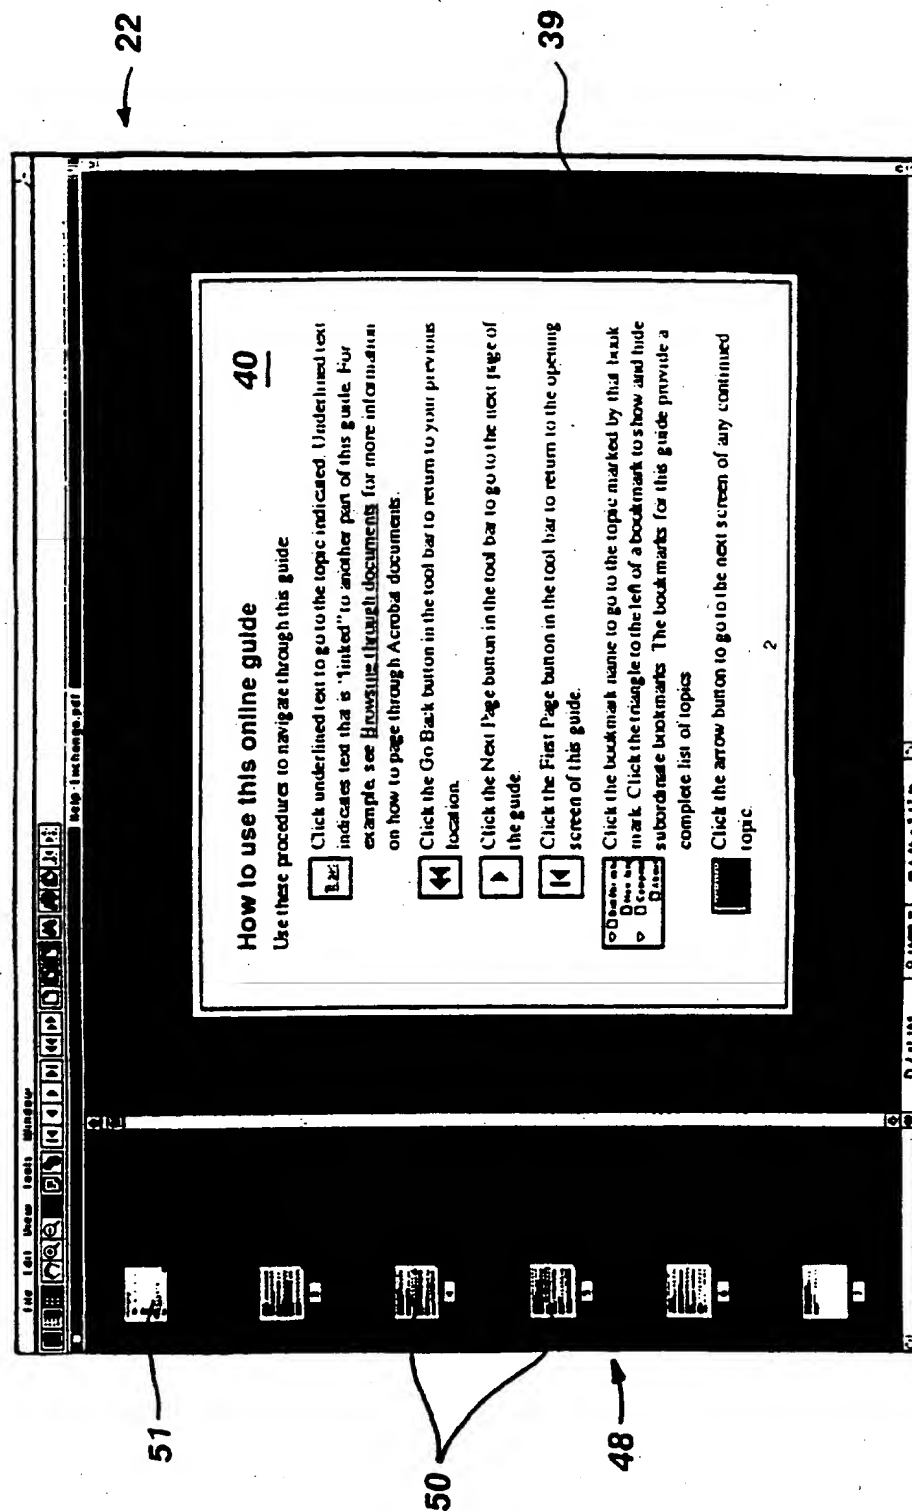


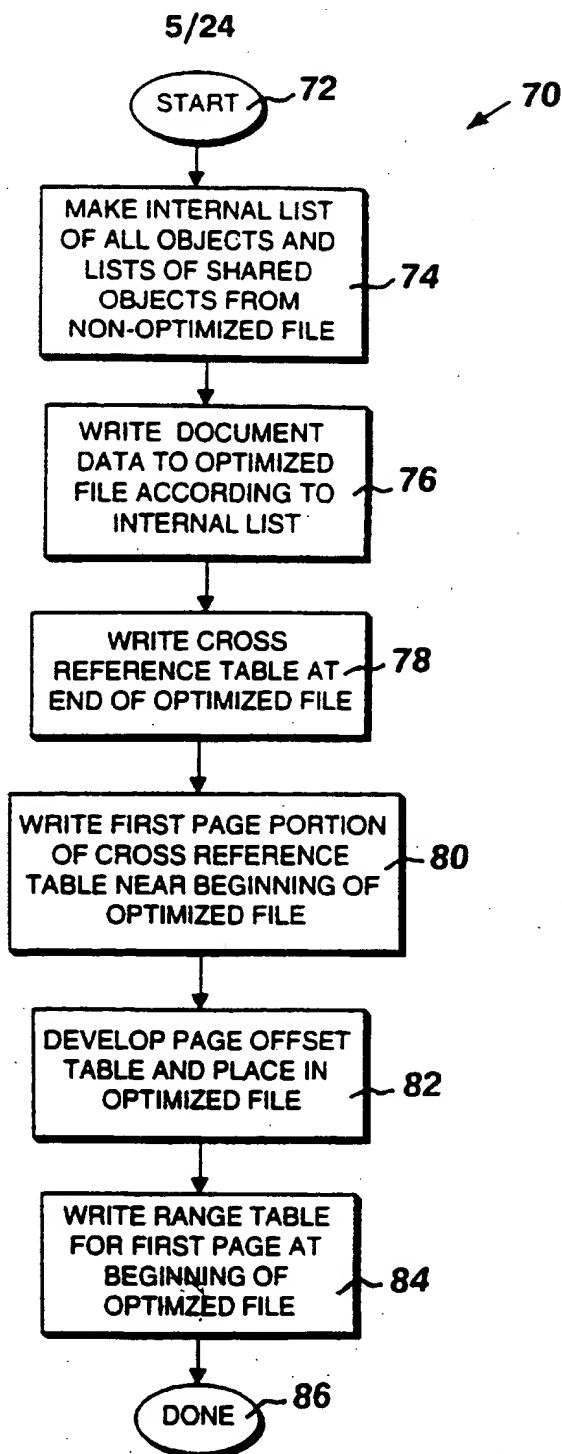
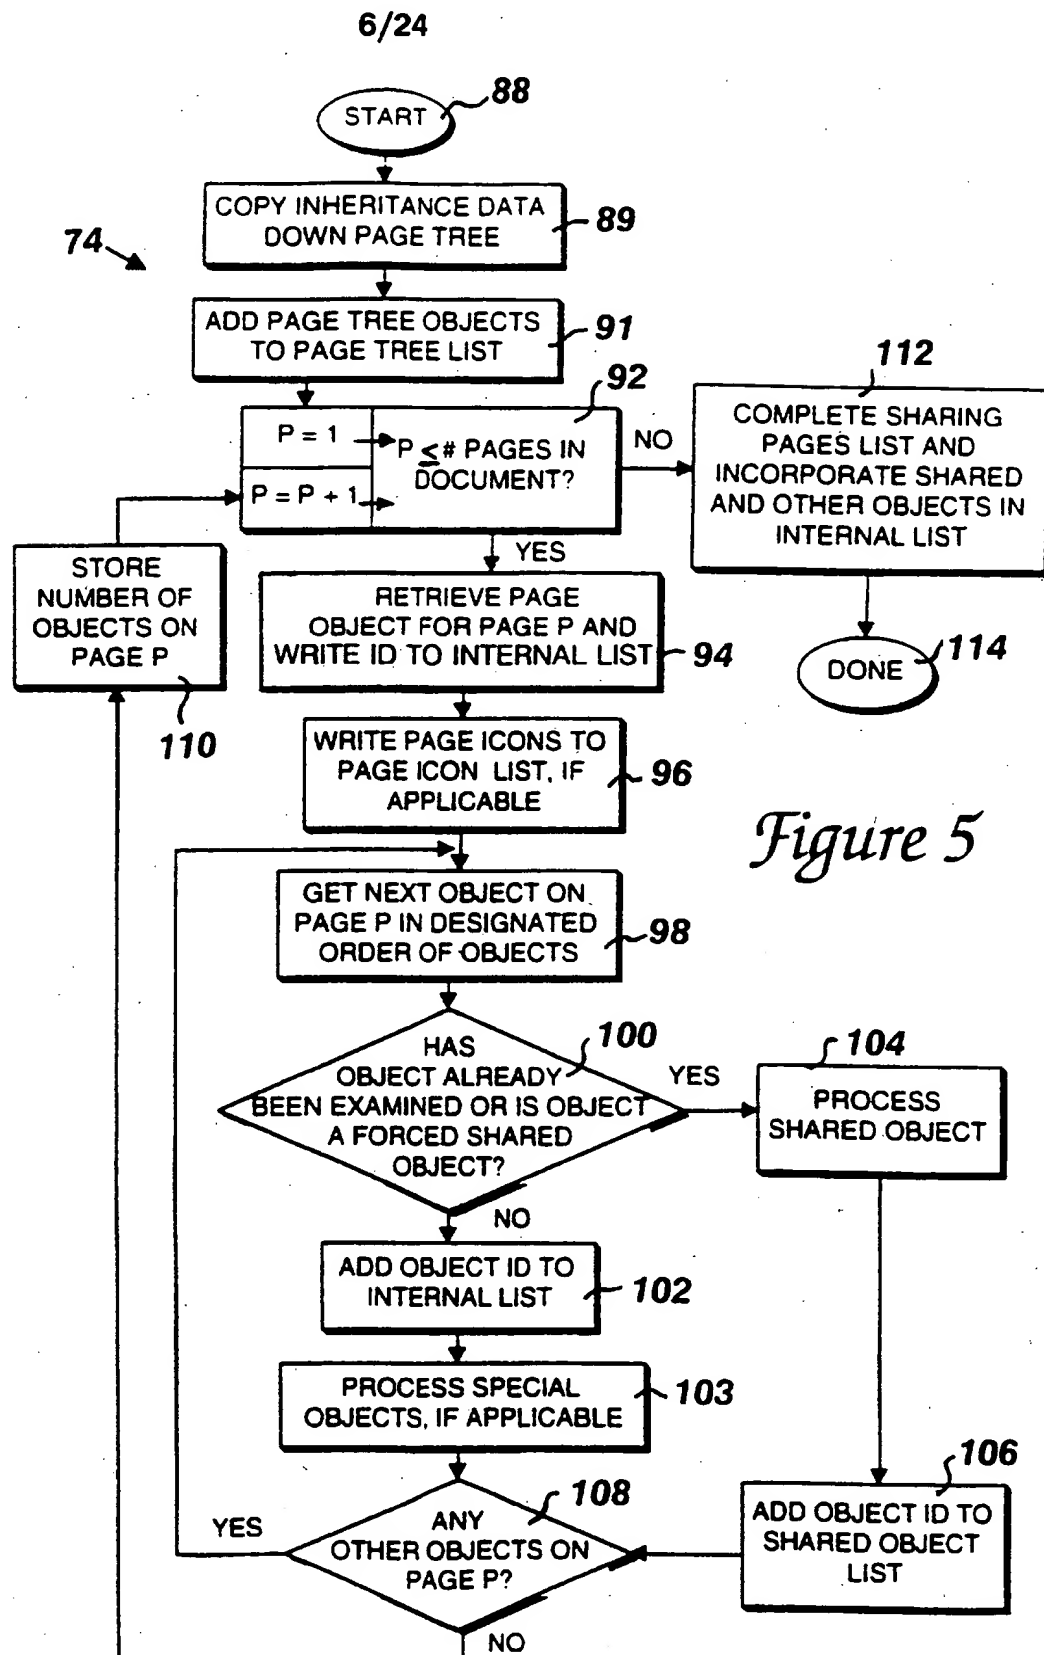


Figure 26



4/24



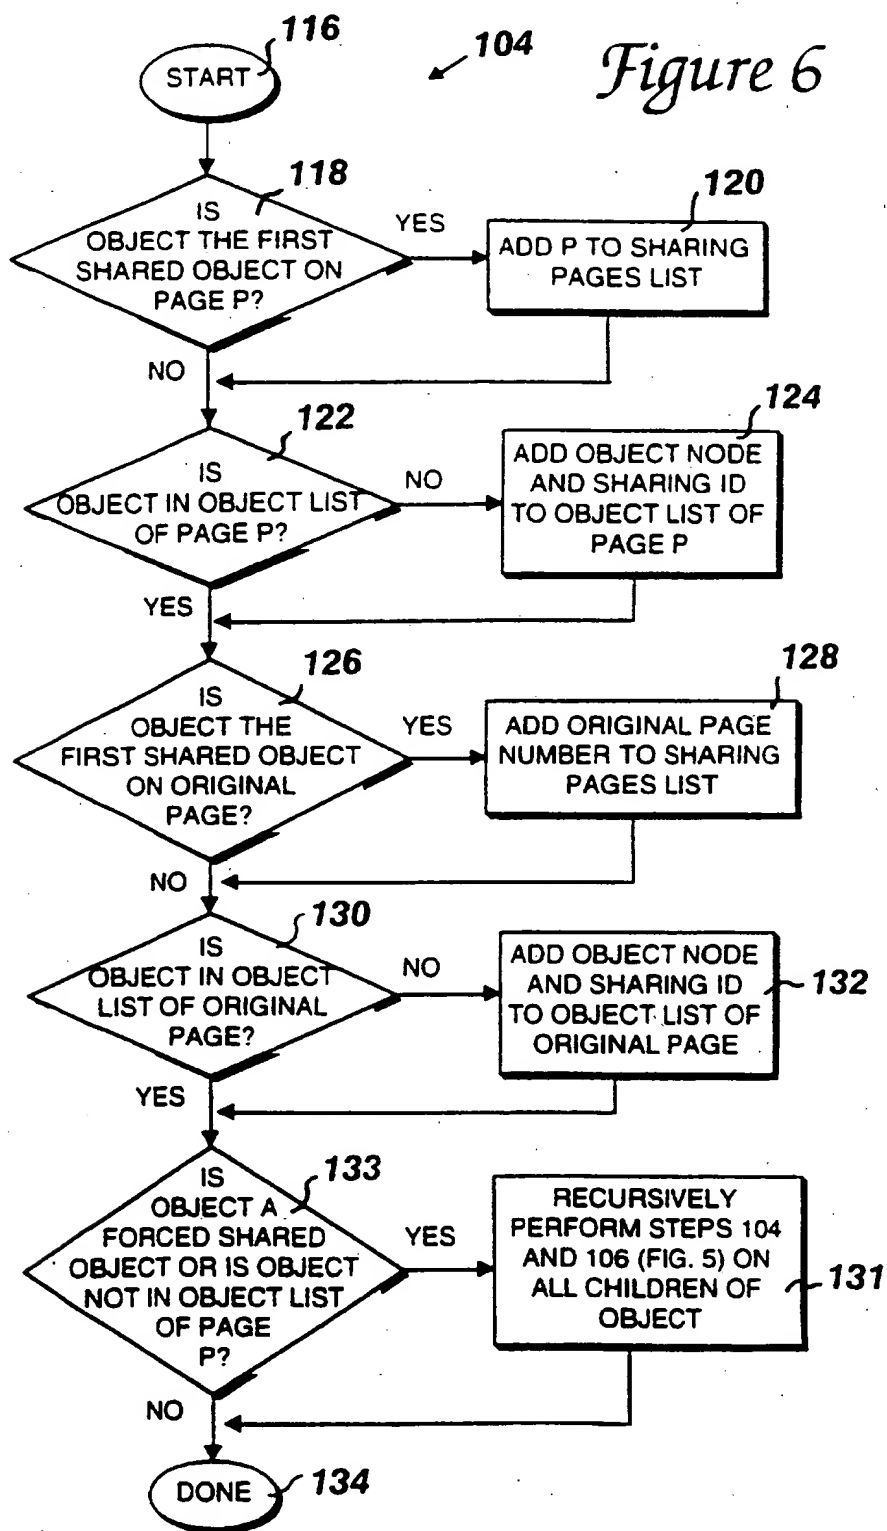
*Figure 4*



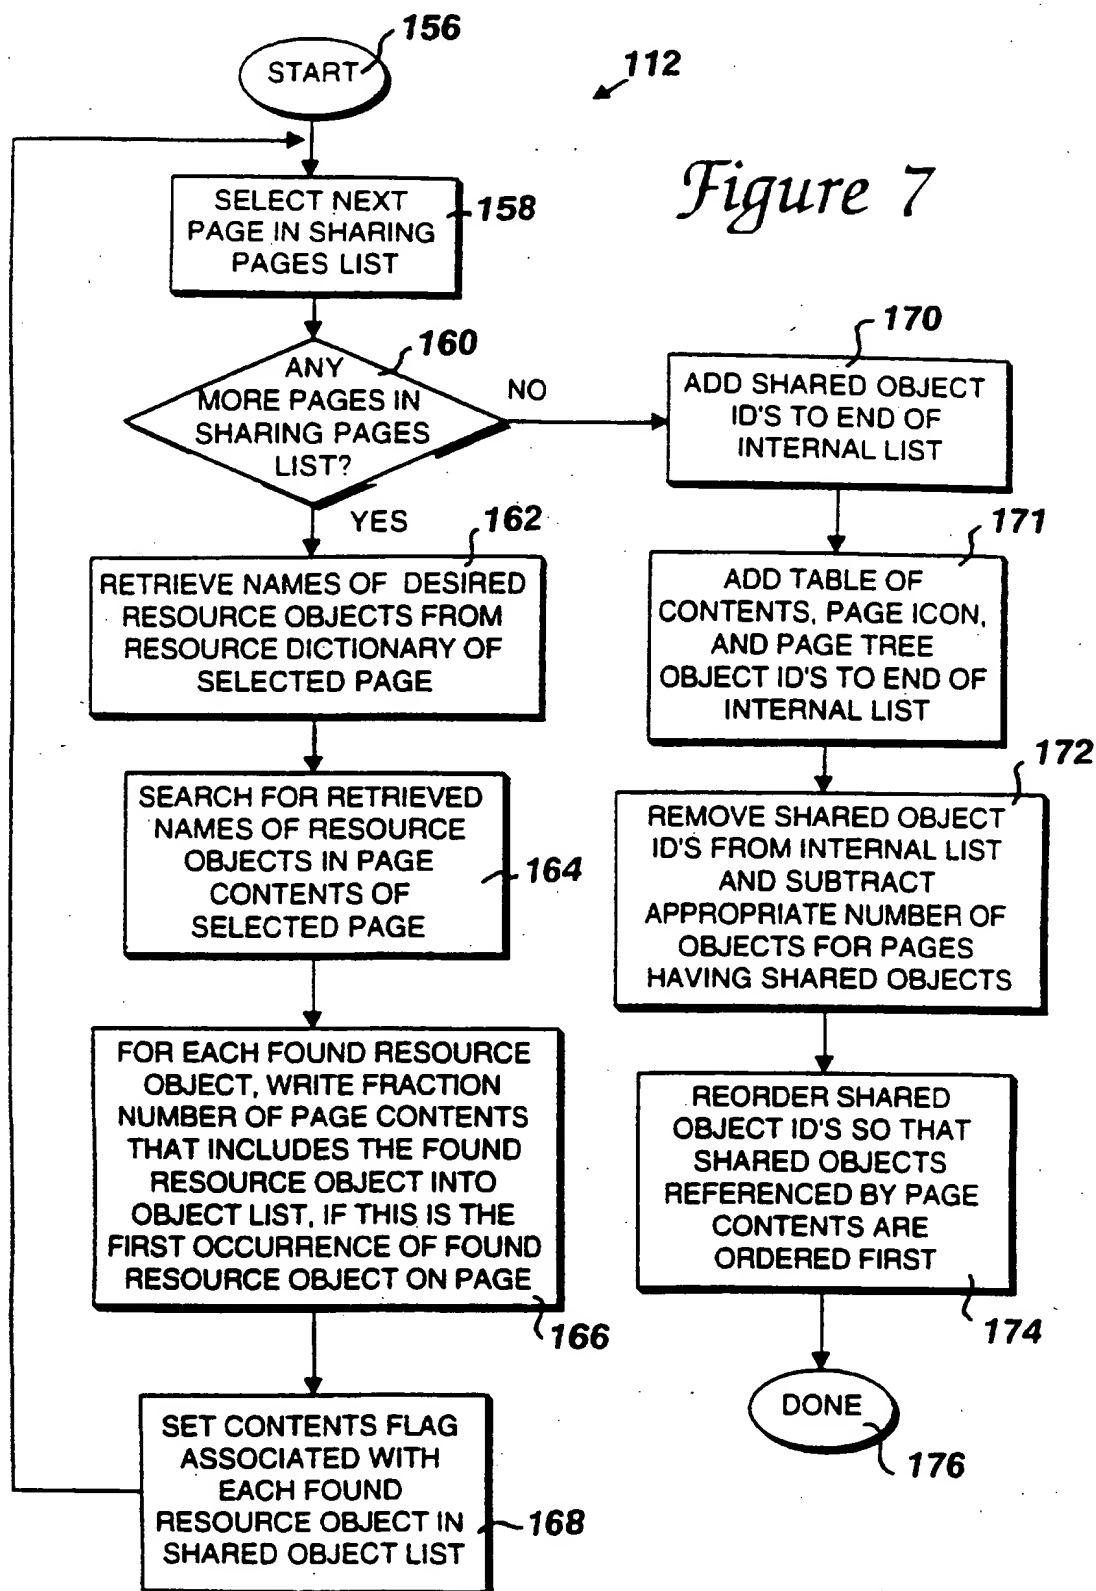
7/24




8/24



9/24



10/24



11/24

68

PAGE OFFSET TABLE

HOUSEKEEPING INFORMATION

NUMBER OF PAGES IN DOCUMENT
 LEAST NUMBER OF OBJECTS ON A PAGE
 LOCATION OF CROSS REFERENCE TABLE
 SIZE (IN BITS) TO REPRESENT # OBJECTS ON PAGE
 LEAST SIZE OF A PAGE IN DOCUMENT
 SIZE (IN BITS) TO REPRESENT LARGEST PAGE LENGTH
 NUMBER OF SHARED OBJECTS IN DOCUMENT
 NUMBER OF NON-CONTENTS SHARED OBJECTS
 LEAST SIZE OF A SHARED OBJECT
 FRACTION SIZE
 SIZE (IN BITS) OF DIVIDEND FOR FRACTION

FOR EACH PAGE

NUMBER OF OBJECTS ON PAGE
 LENGTH (SIZE) OF PAGE
 LENGTHS OF SHARED OBJECTS ON PAGE
 SIGNATURES FOR SHARED OBJECT
 SHARED OBJECT FLAG
 SHARING ID'S
 APPROX. PAGE CONTENTS FRACTION
 FRACTIONS

Figure 8a

RANGE TABLE


66

HOUSEKEEPING INFORMATION

NUMBER OF RANGES
 VERSION
 SIZE (IN BITS) FOR EACH LENGTH
 SIZE (IN BITS) FOR EACH OFFSET
 POINTER TO PAGE OFFSET TABLE
 OFFSETS
 LENGTHS
 SIGNATURE FLAGS
 SIGNATURES

Figure 9a

12/24



13/24






Figure 10


14/24

*Figure 11*

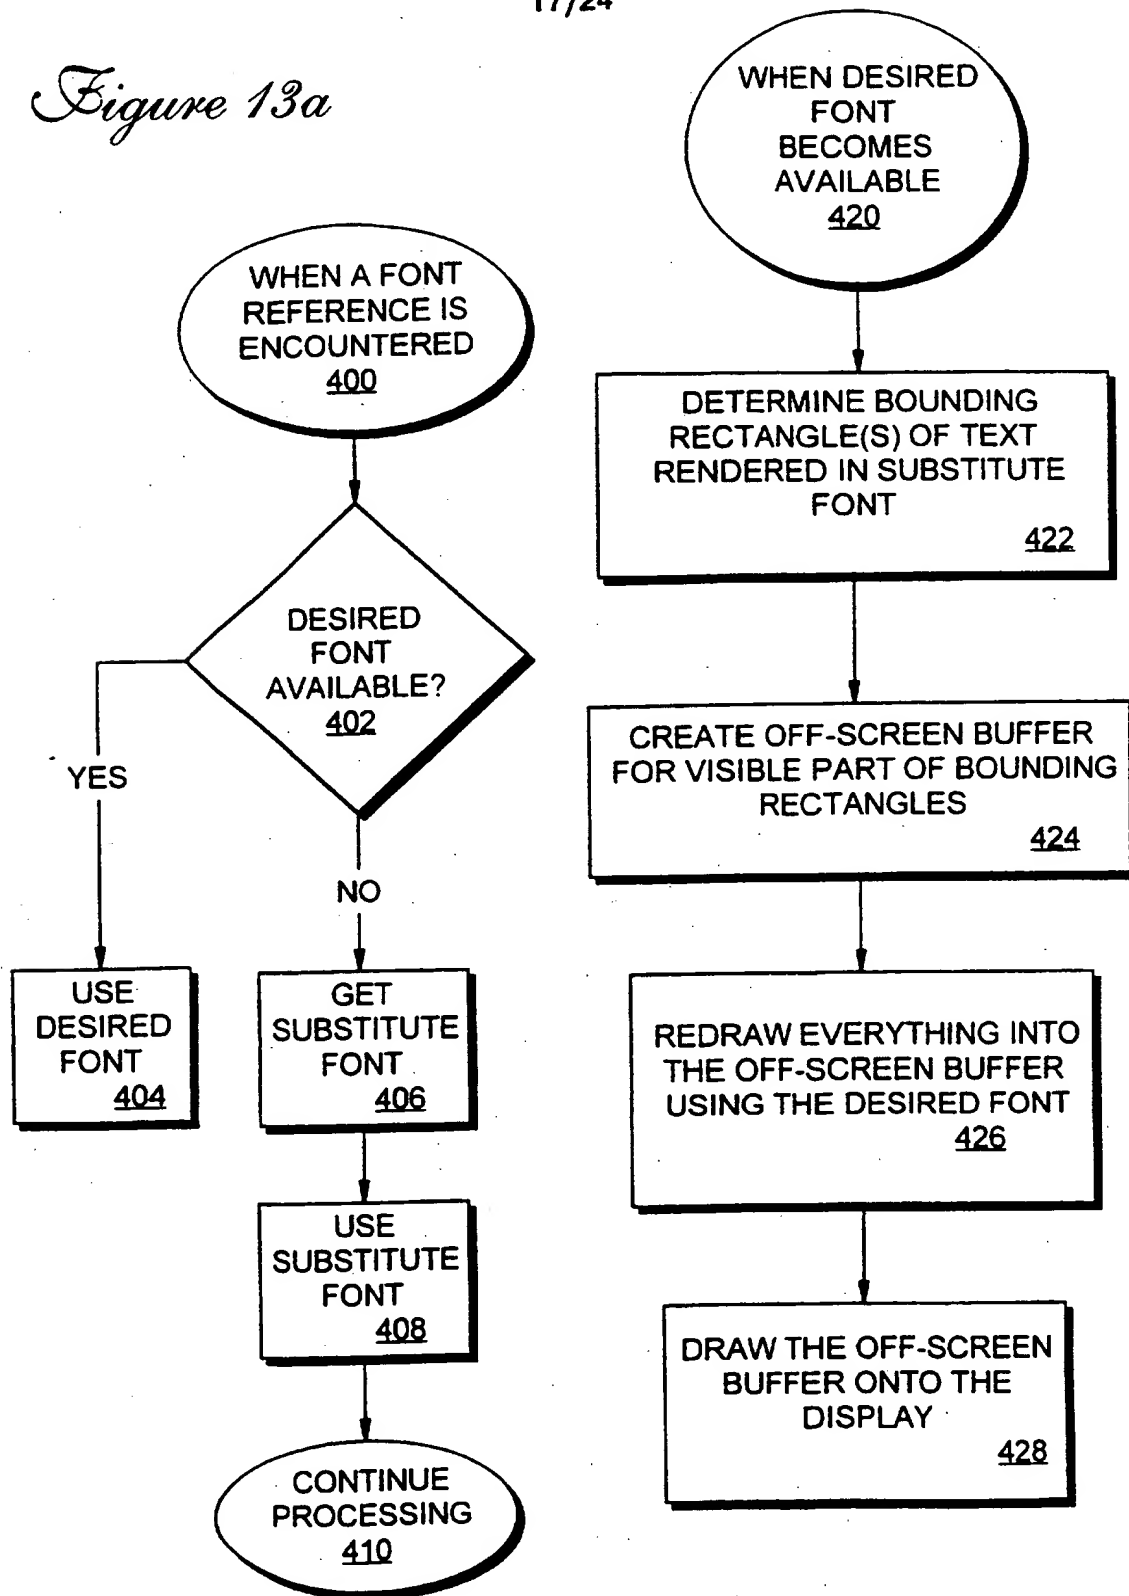
15/24

*Figure 11a*

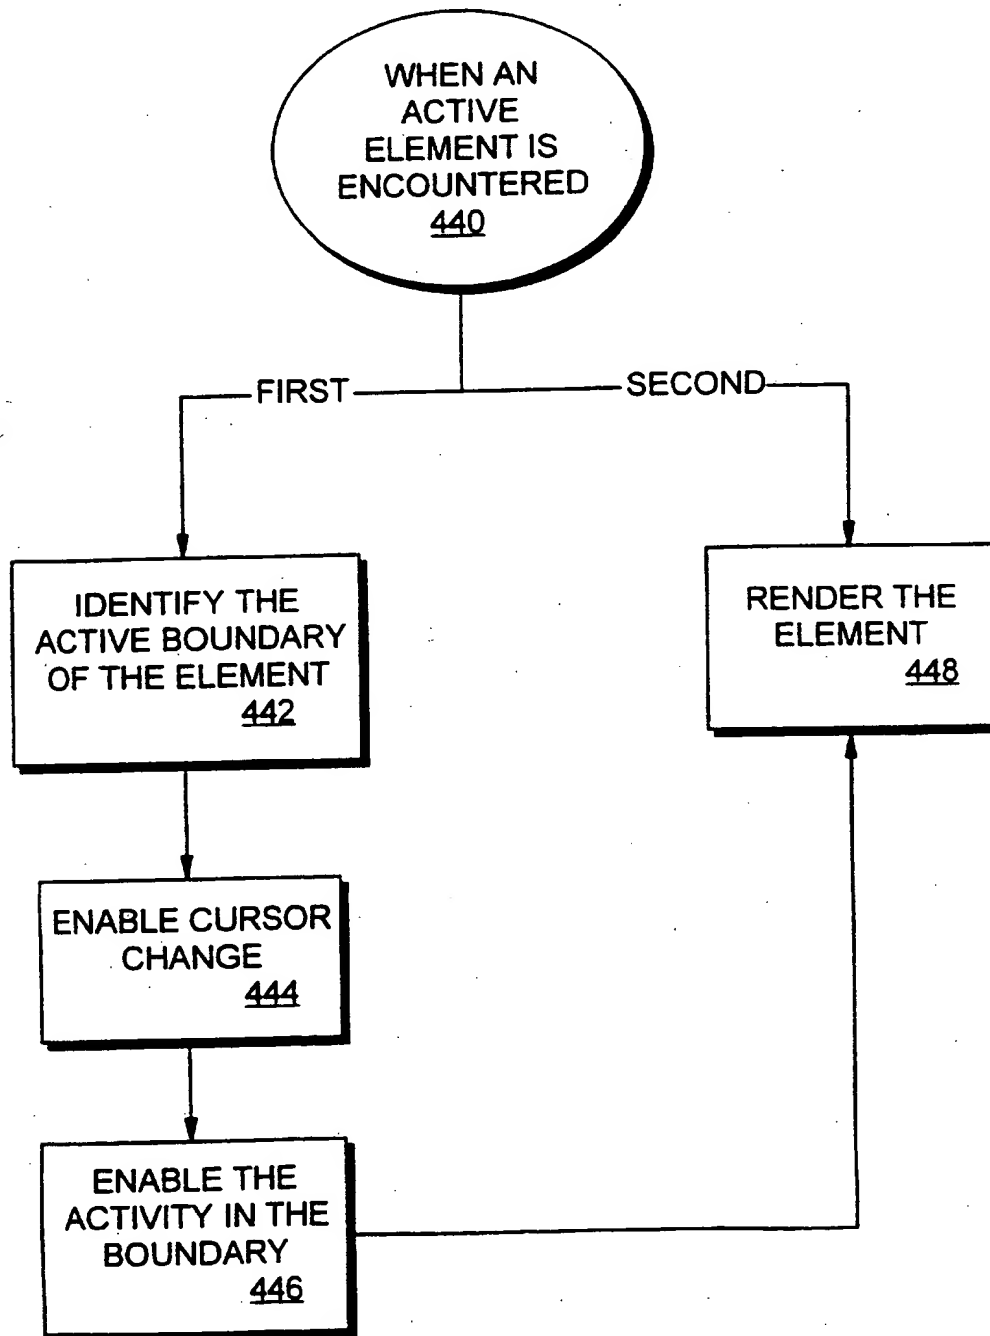
16/24




17/24

Figure 13a

18/24

Figure 13b

19/24

Figure 13c

20/24

Figure 14

HEADER	<u>480</u>
LINEARIZATION PARAMETERS	<u>482</u>
FIRST PAGE CROSS-REFERENCE TABLE	<u>484</u>
PRIMARY HINT TABLE STREAM	<u>486</u>
PAGE OFFSET HINT TABLE	<u>488</u>
SHARED OBJECT HINT TABLE	<u>490</u>
. . .	
FIRST PAGE OBJECTS (INCLUDING SHARED OBJECTS)	<u>492</u>
NEXT PAGE OBJECTS	<u>494</u>
. . .	
LAST PAGE OBJECTS	<u>496</u>
SHARED OBJECTS	<u>498</u>
(OPTIONAL) OVERFLOW HINT TABLE STREAM	<u>500</u>
MAIN CROSS-REFERENCE TABLE	<u>502</u>
TRAILER	<u>504</u>

21/24

Figure 15a

488

PAGE OFFSET HINT TABLE

HEADER INFORMATION

LEAST NUMBER OF OBJECTS IN A PAGE
BITS NEEDED TO REPRESENT NUMBER OF OBJECTS IN PAGE
LEAST LENGTH OF PAGE IN BYTES
BITS NEEDED TO REPRESENT PAGE LENGTH
LEAST START OF CONTENTS OFFSET
BITS NEEDED TO REPRESENT START OF CONTENTS OFFSET
LEAST CONTENTS LENGTH
BITS NEEDED TO REPRESENT CONTENTS LENGTH
BITS NEEDED TO REPRESENT NUMBER OF SHARED OBJECT REFERENCES
BITS NEEDED TO IDENTIFY SHARED OBJECT
BITS NEEDED TO REPRESENT FRACTION
DENOMINATOR USED TO DIVIDE PAGE CONTENTS INTO FRACTIONS

FOR EACH PAGE

NUMBER OF OBJECTS IN PAGE
PAGE LENGTH IN BYTES
NUMBER OF BYTES FROM START OF PAGE TO START OF CONTENTS
STREAM
LENGTH OF CONTENTS IN BYTES
NUMBER OF SHARED OBJECTS REFERENCED FROM PAGE
FOR EACH SHARED OBJECT REFERENCED FROM PAGE:
 SHARED OBJECT IDENTIFIER
 FRACTION GIVING POSITION IN CONTENTS OF FIRST REFERENCE

22/24

*Figure 15b*490

SHARED OBJECT HINT TABLE

HEADER INFORMATION

OBJECT NUMBER OF FIRST OBJECT IN SHARED OBJECTS SECTION
LOCATION OF FIRST OBJECT IN SHARED OBJECTS SECTION
NUMBER OF SHARED OBJECT ENTRIES FOR FIRST PAGE
NUMBER OF SHARED OBJECT ENTRIES FOR SHARED OBJECTS SECTION
BITS NEEDED TO REPRESENT NUMBER OF OBJECTS IN A SHARED
OBJECT GROUP
LEAST LENGTH OF A SHARED OBJECT GROUP IN BYTES
BITS NEEDED TO REPRESENT LENGTH OF A SHARED OBJECT GROUP

FOR EACH SHARED OBJECT GROUP

NUMBER OF OBJECTS IN GROUP
TOTAL LENGTH OF OBJECTS IN GROUP
SIGNATURE PRESENT FLAG
OPTIONAL SIGNATURE

23/24

*Figure 15c*506

THUMBNAIL HINT TABLE

HEADER INFORMATION

OBJECT NUMBER OF FIRST OBJECT IN THUMBNAILS SECTION
LOCATION OF FIRST OBJECT IN THUMBNAILS SECTION
NUMBER OF ENTRIES IN THUMBNAIL TABLE
BITS NEEDED TO REPRESENT COUNT OF THUMBNAIL-LESS PAGES
LEAST LENGTH OF A THUMBNAIL OBJECT IN BYTES
BITS NEEDED TO REPRESENT LENGTH OF A THUMBNAIL
BITS NEEDED TO REPRESENT NUMBER OF SHARED OBJECT REFERENCES

FOR EACH ENTRY

COUNT OF PRECEDING PAGES LACKING THUMBNAILS
LENGTH OF THUMBNAIL OBJECT IN BYTES
NUMBER OF SHARED OBJECTS REFERENCED FROM THUMBNAIL
FOR EACH SHARED OBJECT: SHARED OBJECT IDENTIFIER

24/24

Figure 15d

508

GENERIC HINT TABLE

HEADER INFORMATION

OBJECT NUMBER OF FIRST OBJECT
LOCATION OF FIRST OBJECT
NUMBER OF OBJECT GROUPS
BITS NEEDED TO REPRESENT NUMBER OF OBJECTS IN A GROUP
LEAST LENGTH OF AN OBJECT GROUP IN BYTES
BITS NEEDED TO REPRESENT LENGTH OF AN OBJECT GROUP
BITS NEEDED TO REPRESENT NUMBER OF SHARED OBJECT REFERENCES

FOR EACH OBJECT GROUP

NUMBER OF OBJECTS IN GROUP
TOTAL LENGTH OF OBJECTS IN BYTES
NUMBER OF SHARED OBJECT REFERENCES
FOR EACH SHARED OBJECT REFERENCED FROM GROUP: SHARED OBJECT
IDENTIFIER

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US96/15725

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) : G06F 15/16, 17/21, 13/14

US CL : 395/616, 611, 350, 340, 338

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 395/616, 611, 350, 340, 338

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

None

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

None

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US, A, 4,754,326 (KRAM et al.) 28 June 1988, col.68, line 48- col.74, line 36	1-32
Y	US, A, 4,648,061 (FOSTER) 3 March 1987, col.10, lines 1-67	1-32
A	US, A, 4,532,588 (FOSTER) 30 July 1985	1-32
A,P	US, A, 5,557,722 (DEROSE et al.) 17 September 1996	1-32
A	US, A, 5,040,142 (MORI et al.) 13 August 1991	1-32
A	US, A, 5,315,504 (LEMBLE) 24 May 1994	1-32

☐ Further documents are listed in the continuation of Box C.☐ See patent family annex.

* Special categories of cited documents:

A document defining the general state of the art which is not considered to be of particular relevance

E earlier document published on or after the international filing date

L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

O document referring to an oral disclosure, use, exhibition or other means

P document published prior to the international filing date but later than the priority date claimed

T

later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X

document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y

document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

Z

document member of the same patent family

Date of the actual completion of the international search


03 DECEMBER 1996

Date of mailing of the international search report

11 FEB 1997

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

JEAN R. HOMERE

Telephone No. (703) 308-6647

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US96/15725

Box I Observations where certain claims were found unsearchable (Continuation of item 1 of first sheet)

This international report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. ☐ Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:

3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box II Observations where unity of invention is lacking (Continuation of item 2 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

1. This International Search Authority has found 3 inventions claimed in the International Application covered by the claims indicated below:

This application contains the following inventions or groups of inventions which are not so linked as to form a single inventive concept under PCT Rule 13.1. In order for all inventions to be examined, the appropriate additional examination fees must be paid.

Group I, claim(s) 1-24, drawn to apparatus and methods for using the page offset information to locate the content of the specified page in an electronic document.

Group II, claim(s) 25-29, drawn to methods for controlling display of text and graphical object in an electronic document.

1. ☒ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.

2. ☐ As all searchable claims could be searched without effort justifying an additional fee, this Authority did not invite payment of any additional fee.

3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:

4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest.
☐ No protest accompanied the payment of additional search fees.

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ **BLACK BORDERS**
- ☐ **IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- ☐ **FADED TEXT OR DRAWING**
- ☐ **BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- ☐ **SKEWED/SLANTED IMAGES**
- ☐ **COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- ☐ **GRAY SCALE DOCUMENTS**
- ☐ **LINES OR MARKS ON ORIGINAL DOCUMENT**
- ☐ **REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- ☐ **OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.